Robust Estimation of Private Business Wealth[*]

Job Market Paper

Simon J. Toussaint[†]

November 14, 2024

[Most recent version here]

Abstract

Estimating the market value of private businesses is essential for understanding both aggregate firm dynamics and top wealth inequality, yet these values are inherently unobservable. This paper introduces an econometric approach that treats the gap between true market values and initial estimates as measurement error. I employ time-series restrictions on these errors as moment conditions within a GMM framework, and use the fitted values from these estimations as error-free estimates of private business wealth and capital stocks. Applying this method to Dutch administrative data linking the universe of firms to their owners, I find that aggregate private business wealth increases by 30% of GDP initially, and is more stable than the unadjusted series. Top 1% and 0.1% wealth shares increase by 3–5 percentage points, peaking at 38% and 20%, respectively. Adjusted returns to firm wealth exhibit a steeper gradient across the wealth distribution than unadjusted returns, consistent with models of return heterogeneity.

JEL Classification: D3, E2, G5

Keywords: wealth inequality; firm dynamics; private business wealth; capital stock; return heterogeneity.

[*]I thank Coen Teulings and Bas van Bavel for their guidance and support. I thank my colleagues at Statistics Netherlands for access to and help with the data, especially Arjan Bruil, Jasper Dingerink, and Safan van der Gaauw. I also thank Céline van Essen for help with the data, and Pierre Chaussé for help with the GMM implementation. The paper benefited considerably from helpful comments by Riccardo Cioffi, Wouter den Haan, Albert-Jan Hummel, Tomer Ifergane, Eva Janssens, Xavier Jaravel, Clara Martínez-Toledano, Gabriella Massenz, Ben Moll, Hugo Reichardt, Maarten de Ridder, David Seim, Mara Squicciarini, Daniel Waldenström, and Thomas Winberry, as well as numerous seminar participants.

[†]Utrecht School of Economics and Statistics Netherlands. E-mail: s.j.toussaint@uu.nl

1 Introduction

Private businesses are key drivers of economic activity. In the United States private businesses account for almost half of aggregate sales and profits (Campbell and Robbins 2023). Moreover, their ownership is highly concentrated among the top of the income and wealth distribution (Kopczuk and Zwick 2020). Therefore, to get a good sense of aggregate firm dynamics and top inequality, estimating private business wealth is essential (Saez and Zucman 2016; Smith, Zidar, and Zwick 2023).

The key challenge is that private businesses, by definition, are not listed. Hence, their market values are unobserved. Firm datasets do typically include balance sheet components such as book value of a firm's value and capital. However, these components are unlikely to capture the true market values of a firm. Conceptually, book value of a firm misses intangibles such as goodwill, brand reputation and other things that are capitalized in a market price but not necessarily recorded accurately in a firm balance sheet. Moreover, since these balance sheet components are based on accounting measures, they are inherently backward-looking. In contrast, a firm's market value is forward-looking: it is the expected present value of future income streams.

Most of the literature takes the accounting values at face value, without further adjustment. This risks seriously understating firm value dynamics and top wealth shares. Existing methods to estimate private business wealth rely on strong assumptions, which each risk introducing further measurement error. For instance, a common response is to use industry-level market-to-book ratios, perhaps with some illiquidity discount, and apply these to private firms of similar industries (Bach, Calvet, and Sodini 2020; Damodaran 2012). However, it is highly debatable whether listed firms are representative of their industry, creating selection bias issues. An alternative response is to capitalize cashflows using some estimated discount factor. This is a forward-looking measure and in principle results in correct values; however, this requires the estimation of a (stochastic) discount factor, which opens the door for further measurement errors. A final strategy is to look at resales of private firms, and regress the updated values on firm characteristics (Campbell and Robbins 2023). This approach relies on resales being representative, which is again highly debatable.

The main contribution of this paper is to introduce a new approach that treats the missing market value question as an econometric problem. My strategy is to consider the gap between some initial estimate of market value and true market value as measurement error. Then, we can use tools from the vast literature on measurement error to explicitly account for this measurement error. My approach is structural in nature. From standard neoclassical investment theory, we expect a firm's value to be a linear function of its capital stock; under the neoclassical assumptions, a regression of value on capital has a coefficient equal to Tobin's q. In fact, I show that even in the more general model of Crouzet and Eberly (2023), where firms potentially charge markups and/or face decreasing returns to scale, approximate linearity still holds. Hence, we would expect a firm's market value to be structurally related to its capital stock, and this structural relation should be identifiable using a linear regression. Then, a naive approach would be to take some initial estimate of market value, regress this on capital stock, and keep the fitted values of this regression as the true market values.

Since we do not observe the firm's true value, nor its true capital stock, we cannot directly run this regression. Instead, I use the following three-step approach. First, I construct initial estimates of firms' market value, by capitalizing a three-year moving average of its profits. These initial estimates will be highly contaminated by measurement error, as will the right-hand side of the regression, firms' capital stocks. Hence, in the second step, I regress the estimated market value on the capital stock, using an instrumental variables strategy to filter out the errors. Finally, in an instrumental variables or generalized method of moments (GMM) framework, these instruments then deliver fitted values of capital (from the first stage) and fitted values of market value (from the second stage). Under the identifying assumptions, these fitted values are free from measurement errors. Therefore, I treat these fitted values as the "true" values of market value and capital.

Clearly, my approach relies on valid instruments. I develop an approach using time-series restrictions, exploiting the panel nature of my dataset. As noted by Griliches and Hausman (1986), we can use a variety of estimators (within, first-difference, second-difference, etc.) in panel data settings in the presence of unobserved fixed effects. Under measurement error in the independent variable, these estimators are all biased, but in *different ways*; hence, we can use differences in these estimators as instruments to identify the parameters of the model. This setup yields more instruments than parameters to be estimated. Therefore, we can test over-identifying restrictions, verifying whether the model holds in the data. Specifically, I use an instrument based on the deviations between the within and the first-difference estimator, and instrument based on on the deviations between the first-difference and the second-difference estimator, and an instrument based on the deviations between the within and the second-difference estimator. Under the identifying assumption that the measurement errors are additively separable into a fixed effect and idiosyncratic innovations, these are valid instruments. In extensions, I weaken these assumptions to allow the innovations in measurement error to follow a MA(1) process, and derive valid instruments under this assumption.

I apply my method to administrative data from the Netherlands 2008–2020. I link the universe of incorporated firms to their owners, which enables me to investigate dynamics both at the firm and the household level. For my econometric methods, I look at the firm data only. I use the universe of corporate income tax returns, which have rich information on firms' balance sheet components and income statements. I construct the capital stock as the sum of a firm's physical and intangible capital. I construct my raw estimate of a firm's market value by taking a three-year moving average of its profits as my forecast of its future profits, capitalized by three different choices of discount factor. Obviously, these estimates are all imperfect and subject to error. This is why my GMM approach is important: it filters out the error in these estimations by retaining only the components that are structurally related to a firm's true capital stock. GMM estimates across these different initial estimates are quantitavely very similar, lending confidence to the validity of the instrumental-variable procedure.

My first empirical application is to investigate aggregate firm dynamics. My adjustments increase the total value of private business wealth markedly, in particular in the early years of my sample. The unadjusted microdata show a steep increase in this wealth component's aggregate value from €200 billion in 2008 to around €450 billion a decade later. In contrast, the updated values are almost €150 billion are higher in early years, an increase of around 30% of GDP. The adjusted series stays more stable over the period. Hence, at the end of the sample, the microdata and the updated series almost coincide. This shows that while the procedure does generally increase the measured values, it does not do so mechanically. This also addresses concerns voiced by Cochrane (2020), Greenwald, Leombroni, Lustig, and Van Nieuwerburgh (2021), and Fagereng et al. (2024) and others whether the increase in wealth inequality is merely a valuation effect ("paper gains"). My results show that it is the combination of discount rates and profits that drive the changes in private business wealth, not merely declining discount rates.

I analyze why the book values are lower than my updated series and conclude that this is likely driven by portfolio reallocations of investors as interest rates declined. The Dutch tax system taxes savings and directly-held financial assets quite heavily, effectively using a wealth tax. The decline in deposit interest rates since 2013 coincided with a major reallocation of capital towards private firms, which are are not subject to the wealth tax.[1] This suggests that book values of Dutch firms were quite heavily undervalued at the beginning of my period, and mechanically increased in value as capital started flowing in. In contrast, the firms' capital stock remained quite stable, which means that my updated measures also remained stable. I conclude that while book-value measures of private firm wealth are vulnerable to fiscal manipulation, my alternative series is more robust, as well as measuring the economically relevant concept.

I also use the aggregate data to investigate why the coefficients from my GMM procedure, which range between 3.5 and 4.5 on average, are so sizable. Intuition based on Tobin's q suggests that these values should be lower, on the order of 1.5.[2] However, the model developed in Section 2 suggests two answers to this question. First, if either market power and/or decreasing returns to scale are present, firms have positive rents in equilibrium. The net present value of these rents shows up in the regression coefficient. I estimate the structural parameters of the model and conclude that rents are indeed present. In fact, rents are substantially larger than the estimates in Crouzet and Eberly (2023) for U.S. listed corporations. This suggests that

1. Dividends and realized capital gains are subject to income taxation. See Section 4 for an extensive discussion of the Dutch capital income taxation system.

2. According to the United States Flow of Funds, Tobin's average Q averages around 1.3 since 2008 in the United States.

applying the model to the universe of firms reveals larger rents than just among listed firms. However, applying the structural model also reveals puzzlingly low estimated capital costs as a share of total firm costs. This suggests that the capital stock is underrecorded in firms' accounts. Upwardly adjusting capital stocks would decrease the other parameters of the model, such that the fitted values would remain the same.

My second empirical application concerns wealth inequality. My results have major implications for top wealth shares, since private firms are extremely concentrated at the top of the wealth distribution: it accounts for almost 80% of the top 0.01%'s portfolio. Based on the lower range of the fitted values, the top 1% wealth share increases by 3 percentage points on average relative to the headline series, to about 36% at its peak. The top 0.1% share increases by a similar amount, to about 17%. These are the lower-range results from my specifications. If I use the upper range, the top 1% share increases further to almost 39% at its peak, and the top 0.1% share almost equals 20%. Clearly, these are estimates, and the exact quantitative upward adjustment is open to reasonable debate. Nevertheless, my results indicate unambiguously that inequality is underestimated in the microdata. Based on these adjustments, the Netherlands would rank among the most unequal developed countries in the world on wealth inequality. However, series from other countries likewise suffer from underestimating private business wealth (Kopczuk and Zwick 2020); hence it seems plausible that my procedure would result in higher top wealth shares in other countries as well.

A final implication of my results concerns the debate on return heterogeneity. As observed by Gabaix, Lasry, Lions, and Moll (2016), existing models of wealth inequality have difficulty accounting for the speed of inequality increases. Models with heterogeneous returns, however, are able to account for these dynamics. Several papers have documented return heterogeneity (Fagereng, Guiso, Malacrino, and Pistaferri 2020; Bach, Calvet, and Sodini 2020; Xavier 2020), which seems to confirm the hypothesis. However, I show that these results are highly sensitive to measurement error. Since returns have wealth (or assets) in the denominator, a regression of returns on wealth yields mechanical correlations between left- and right-hand side if there is measurement error in wealth. The sign of this bias depends on whether the numerator is also measured with error, as seems likely in the case of private business wealth. These theoretical results underscore the need to take measurement seriously in accounting for wealth dynamics. I find that adjusting firms' values makes a major difference. Returns based on unadjusted returns are essentially flat or weakly increasing across the firm size distribution. Adjusted returns, on the other hand, strongly increase with firm size, with the gradient becoming the steepest in the top decile. These conclusions also hold when we aggregate to the household level. Therefore, heterogeneity in returns (either at the firm or firm-owner level) is substantial, with most of the heterogeneity coming from the top. Unadjusted returns, due to measurement error or other noise, understate this effect. My results also help reconcile an apparent contradiction in the literature. Most papers focusing on economic returns find that returns increase with wealth. In contrast, a recent paper by Boar, Gorea, and Midrigan (2021) finds that *accounting* returns are decreasing in firm size. They interpret this finding with an entrepreneurship model featuring financing frictions. My results show that both are simultaneously true: accounting returns are flat or even decreasing in the firm size distribution, while economic returns are increasing in firm size.

Related Literature: This paper integrates several strands of literature. Methodologically, I build on a long literature seeking to identify models in the presence of measurement error, as reviewed by Schennach (2016, 2022). Within this literature, my approach builds on Griliches and Hausman (1986), who study measurement error in panel data settings. I adapt their approach to the study of private business valuation, which to the best of my knowledge has not been done before. This question fits with the influential stream of papers in corporate finance and macroeconomics which seek to robustly test the neoclassical investment model, as developed by Hayashi (1982) and Abel and Eberly (1994) and many others. These papers typically use higher-order moment restrictions in a GMM procedure to identify investment-q regressions (Erickson and Whited 2000, 2012; Erickson, Jiang, and Whited 2014), which goes back to insights by Reiersøl (1950), Kapteyn and Wansbeek (1983) and Lewbel, Schennach, and Zhang (2023) and many others. Besides using the Griliches-Hausman framework instead of higher-order moments, my approach differs from the existing literature in treating private business wealth as the fitted values of a regression. In contrast, most papers in this literature are more interested in estimation of the coefficient in investment-q regressions, and therefore do not focus on using the fitted values as error-free estimates of (private) business wealth.

Substantively, my paper relates to many papers in finance that seek to account for the correct valuation of private businesses (Damodaran 2012; Kaplan and Schoar 2005; Korteweg 2019; Gupta and Van Nieuwerburgh 2021). More broadly, my methods and contributions fit within the literature in empirical corporate finance seeking to robustly estimate key firm objects, such as intangibles (Crouzet and Eberly 2021, 2023), labor and capital shares (Barkai 2020; Karabarbounis and Neiman 2014), and firm-level discount rates (Gormsen and Huber 2023, 2024). Central to all these literatures is a concern with correctly measuring firm capital and risk (and hence value), which is the central goal of this paper. Novel to my approach is the explicit econometric approach, where most existing methods either use an asset-pricing model (e.g., Gupta and Van Nieuwerburgh 2021), or advocate empirical proxies such as applying industry-level market-book ratios from listed firms (Damodaran 2012).

My results have major implications for the wealth inequality literature (Saez and Zucman 2016; Smith, Zidar, and Zwick 2023; Saez and Zucman 2020). Many authors have noted concerns with the reliable measurement of top wealth shares in the presence of unlisted private businesses (Smith, Zidar, and Zwick 2023; Kopczuk and Zwick 2020; Toussaint, van Bavel, Salverda, and Teulings 2020; Toussaint, de Vicq, Moatsos, and van der Valk 2022). As yet, there is no agreed-upon method to address these concerns. My method is straightforward to implement for researchers with access to firm balance sheet data, and thus provides a theoretically grounded, econometrically robust first step toward measuring the true extent of top wealth inequality.

Finally, my results relate to the debate on return heterogeneity (Gabaix, Lasry, Lions, and Moll 2016; Fagereng, Guiso, Malacrino, and Pistaferri 2020; Bach, Calvet, and Sodini 2020; Xavier 2020). Measurements from Norway and Sweden seem to establish correlations between wealth and returns, but it is still unclear what the underlying mechanism is underlying these features. Explanations include Jones and Kim (2018), Kacperczyk, Nosal, and Stevens (2019), Gerritsen, Jacobs, Spiritus, and Rusu (2024), and Guvenen et al. (2023). My results first clarify that indeed, returns are heterogeneous, and that most of the action is concentrated at the top of the firm and firm-owner distributions. Future theoretical and empirical work should therefore focus on the interplay between high-return firms and high-return owners. In addition, my results show that there is a significant difference between heterogeneity in accounting returns and economic returns; this calls for further work along the lines of Boar, Gorea, and Midrigan (2021) to study differences between the two.

Paper Outline: The rest of the paper is organized as follows. Section 2 discusses the structural model that underpins the use of regressions of firm value on capital. In Section 3, I discuss the two estimation strategies I employ to overcome measurement error. Section 4 introduces the data used and provides summary statistics. Section 5 shows the results of the econometric strategies. The next three sections discuss several implications. Aggregate levels of private business wealth are discussed in Section 6, distributional implications in Section 7, and the consequences for heterogeneous returns are covered in Section 8. Section 9 concludes.

2 Framework

In neoclassical models of investment, market value is related to capital via Tobin's marginal q, the shadow value of an additional unit of investment. Hayashi (1982) shows that if firms are price takers and face constant returns to scale, marginal q and average Q (the firm's ratio of market value to capital stock) are the same: Market value is a linear function of capital, or

$$V_t = qK_t. \tag{1}$$

This result rests on strong assumptions. Besides the Walrasian assumptions already mentioned, the firm is assumed to be able to adjust capital subject to a convex adjustment cost function $\Phi(K_t)$. The convexity ensures that the firm smoothly and continually adjusts its investment; irreversibilities are ruled out. However, Abel and Eberly (1994) show that mild forms of fixed costs in investment can also be accommodated in this framework. As long as the fixed costs are proportional to the profit function Π_t, firm value can still be written as a linear function of capital.

Under these assumptions, my strategy is straightforward: compute some measure of market value (which may be ridden with measurement error), regress it on capital (instrumented by some instrument to correct for endogeneity), and use the fitted values as the true measure of market value. However, we need to verify whether this argument also holds in more general models where firms have market power and face decreasing returns to scale. I now construct a model along those lines and show that, approximately, linearity holds.

The Model: The model is based on Crouzet and Eberly (2023). Time is discrete. I use capital letters for aggregate variables, and lower-case letters for firm-level or input-level variables. Consider a representative monopolistic firm that uses G distinct variable inputs $\{m_{gt}\}_{g=1}^{G}$, as well as capital K, to produce and sell output to consumers, whose demand function is given by:

$$C_t = P_t^{-\frac{\lambda}{\lambda-1}} D_t,$$

where C_t is consumption of the final good, P_t is the price, D_t indexes aggregate demand, and $\lambda \geq 1$ is the firm's markup over marginal cost. The price of variable input g is given by p_t^g. The total input of capital K_t (which is made up of tangible and intangible capital), is quasi-fixed: it is chosen dynamically but cannot be modified immediately. The firm has a production function

$$F(A_t, K_t, M_t) = A_t \left(K_t^{\delta} M_t^{1-\delta}\right)^{\eta}, \tag{2}$$

where A_t is an aggregate productivity shifter, and M is a Cobb-Douglas aggregator of variable inputs, including labor:

$$M := \prod_{g=1}^{G} m_{gt}^{\nu_g}, \quad \sum_{g=1}^{G} \nu_g = 1. \tag{3}$$

The firm can have both rents from market power on the product power, given by λ, and quasi-rents from decreasing returns to scale of degree η in capital and variable inputs. To obtain an interior solution, I follow Crouzet and Eberly (2023) and assume $\eta \leq \lambda$. Interesting cases are typically those where $\eta \leq 1$, so that the firm has decreasing returns to scale; but constant or increasing returns are also covered in this model as long as the markup λ is sufficiently high.

The firm chooses variable inputs and an output price to maximize its profits:

$$\Pi_t = \max_{\{m_{gt}\}_{g=1}^{G}, P_t} P_t^{-\frac{\lambda}{\lambda-1}} D_t - \sum_{g=1}^{G} p_{gt} m_{gt}, \tag{4}$$

$$\text{s.t.} \quad A_t \left(K_t^{\delta} M_t^{1-\delta}\right)^{\eta} \geq P_t^{-\frac{\lambda}{\lambda-1}} D_t. \tag{5}$$

An attractive property of this model is that the optimization problem is static; this is done by not explicitly modeling capital adjustment. Doing so, for instance by including a convex adjustment cost function $\Phi(K_t)$, would complicate the algebra without altering the substantive results. A second advantage of the model is that it admits a closed-form representation of the profit function Π_t that is highly tractable, as derived by Crouzet and Eberly (2023) and reproduced in Appendix A for completeness:

$$\Pi_t = H_t^{1-\frac{1}{\mu}} K_t^{\frac{1}{\mu}} \tag{6}$$

where:

$$\mu := 1 + \frac{\varphi - 1}{\delta} \geq 1, \tag{7}$$

$$\varphi := \frac{\lambda}{\eta} \geq 1, \tag{8}$$

and H_t is a combination of parameters and variables that does not depend on K_t:

$$H_t := \left(\frac{\varphi}{1-\delta}\right)^{-\frac{\varphi}{\varphi-1}} \left(\frac{\varphi}{1-\delta} - 1\right)^{\frac{\varphi-(1-\delta)}{\varphi-1}} D_t^{\frac{\varphi-\eta}{\varphi-1}} \mathcal{P}_t^{-\frac{1-\delta}{\varphi-1}} A_t^{\frac{1}{\eta(\varphi-1)}}$$

where $\mathcal{P}_t$ is a weighted aggregate of input prices p_{gt}:

$$\mathcal{P}_t := \prod_{g=1}^{G} \left(\frac{p_{gt}}{\nu_g}\right)^{\nu_g}.$$

This result has several implications. First, the impact of market power and returns to scale is pinned down by the reduced-form parameter $\varphi = \lambda/\eta$ (although η and λ are generally not separately identified). If the firm has constant returns to scale ($\eta = 1$) and is a price taker ($\lambda = 1$), $\mu = 1$ and profits are a linear function of capital.[3] In the general case where $\varphi > 1$, equation (7) shows that this impacts the profit curvature parameter μ, which will then be larger than 1. I now show that even when $\mu > 1$, profits are still approximately linear in capital:

$$\begin{aligned}\Pi_t &\approx \sum_{n=0}^{\infty} (K_t - 1)^n \binom{\frac{1}{\mu}}{n} H_t^{\frac{\mu-1}{\mu}} \\ &\approx \underbrace{H_t^{\frac{\mu-1}{\mu}} \left(1 - \frac{1}{\mu}\right)}_{=:\alpha} + \underbrace{\frac{1}{\mu} H_t^{\frac{\mu-1}{\mu}}}_{=:\beta} K_t + \frac{\frac{1}{\mu} - 1}{2\mu} (K_t - 1)^2 H_t^{\frac{\mu-1}{\mu}} + O\left(K^3\right).\end{aligned} \tag{9}$$

Equation (9) shows that, to first order, profits are linear in capital. The second-order term vanishes for minor deviations from $\mu = 1$. Crouzet and Eberly (2023) find in their application of American listed corporations a μ ranging from 0.984 to 1.290.[4] These values imply that the fraction premultiplying the second-order term ranges from 0.01 to −0.09, which are one to two orders of magnitude smaller than β. Hence, I conclude that even with market power and returns to scale, approximate linearity holds. As a result, firm value, which is the discounted present value of profits, is also a linear function of capital:

$$V_t = \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho}\right)^t \Pi_t = \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho}\right)^t (\alpha + \beta K_t). \tag{10}$$

Since the firm has rents, its value will include the capitalized value of expected future rents. Through Equation (6), the profit function includes the capitalized value of rents through H_t, and the importance of rents is governed by μ. Under the

3. In fact, in this static framework, profits are identically equal to capital when $\mu = 1$. In a more general formulation with convex adjustment costs to capital, this coefficient (which equals marginal q) would be larger than one (Hayashi 1982).

4. Values smaller than 1 are strictly speaking not possible under the model written above; however, it is clear that this value would be statistically hard to tell apart from 1, although Crouzet and Eberly do not report standard errors.

frictionless benchmark ($\mu = 1$), rents do not show up in the profit function and hence also not in firm value.

So far, the model has been set up without uncertainty. Incorporating this into the model is straightforward, by letting $\{\Pi_t\}$ follow a stochastic process. In particular, to facilitate the operationalization of my model, I make an important assumption:

Assumption 1. *The profit function* $\{\Pi_t\}$ *follows a martingale process.*

Assumption 1 drastically simplifies the computation of the expected present value (equation (10)), since it implies that we can use the current value of a firm's profits as the best linear forecast of future profits. Since profits tend to be highly volatile, I operationalize this assumption by taking a three-year moving average of firm i's profits, $\overline{\pi}_{it}$ (lower-case letters denoting firm-level variables). Then, we estimate v_{it} as

$$v_{it} = \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho}\right)^t \overline{\pi}_{it} = \frac{1+\rho}{\rho}\overline{\pi}_{it}. \tag{11}$$

Estimating firm value in this way is grounded in the model presented above (and holds in most economic models of firm investment, see Crouzet and Eberly (2023)). However, this measure is highly sensitive to the stochastic discount rate ρ, which is an object to be estimated. Even absent measurement error, the fact that we use estimated values for ρ will introduce estimation uncertainty into v_{it}, biasing the results. Moreover, it is quite plausible that the discount rate and profits are measured with error, further attenuating the results.

My estimation strategy seeks to overcome this measurement error issue. It is based on the observation, from equation (9), that firm value is an approximately linear function of the capital stock. However, the capital stock is also likely to be mismeasured, since intangibles are poorly recorded in a firm's accounts. Importantly, mismeasured intangibles and other mismeasurements of the capital stock would affect both the capital stock and firm profits, hence my estimate of firm value. As I show in Section 3, this introduces endogeneity; moreover, left- and right-hand side are mechanically correlated. To overcome this endogeneity, we need some valid instrument that is correlated with the true capital stock, but uncorrelated with the measurement errors. I detail in the next section how I construct this instrument. Summarizing, my procedure to estimate firms' market values is as follows:

1. Construct an estimate of v_{it}, using some discount rate ρ and profits $\overline{\pi}_{it}$. Call this estimate y_{it}. Take a measure of the firm's capital stock from its accounts; call this measure x_{it}.

2. Construct a valid instrument z_{it} that is correlated with the true capital stock k_{it} but uncorrelated with the measurement error.

3. Regress y_{it} on x_{it}, instrumented with z_{it}; obtain fitted values $\widehat{y}$, which are free from measurement error.

In the next section, I more elaborately formulate the econometric problem, and discuss my identification strategy.

3 Estimation Strategy

3.1 Measurement Error

I write the measurement error problem as follows. We are interested in the model

$$v_{it} = \alpha + k_{it}\beta + \varepsilon_{it} \tag{12}$$

where v_{it} is the market value of firm i in year t, and k is the firm's capital stock. ε_{it} is an error term capturing all other factors that might influence a firm's market value. Under the frictionless conditions of Hayashi (1982), $\beta = q$, i.e., β equals Tobin's marginal q, the shadow value of an additional unit of investment. For my purposes, it is irrelevant whether $\beta = q$, as long as a linear and positive structural relation between capital and value holds. Equation (9) shows that this approximately holds even in models with markups and decreasing returns to scale.

Instead of observing a firm's true market value and capital stock, however, we observe the accounting values of these variables, which we treat as measurement-error-contaminated proxies for the true variables. For simplicity, I assume that firm value and capital are contaminated with the same measurement error ψ_{it}, which contains a fixed firm-specific component χ_i and idiosyncratic innovations ξ_{it}. In the context of this paper, we can think of χ_i as the initial or average error when first writing down the firm's capital stock in the accounts, while ξ_{it} are the noise in the annual investments (which cumulate into the capital stock). Note that all results go through if we allow for different measurement erors in v and k, as long as the identifying assumptions apply to them symmetrically.[5] Inserting $\psi_{it} := \chi_i + \xi_{it}$ into the structural model, we have:

$$x_{it} = k_{it} + \chi_i + \xi_{it} \tag{13}$$

$$y_{it} = v_{it} + \chi_i + \xi_{it}. \tag{14}$$

We can rewrite this system to obtain the reduced form

$$y_{it} = x_{it}\beta + \underbrace{\alpha + (1-\beta)\chi_i}_{=:\gamma_i} + \underbrace{\varepsilon_{it} + (1-\beta)\xi_{it}}_{=:u_{it}}. \tag{15}$$

A regression of y_{it} on x_{it} is biased for two reasons: first, because of the fixed effect γ_i, which absorbs the intercept α, and second, because the error term u_{it} is correlated with the regressor x_{it}. Stacking across observations, we have

$$\boldsymbol{y} = \boldsymbol{x}\beta + \boldsymbol{\gamma} + \boldsymbol{u} = \boldsymbol{x}\beta + (1-\beta)\boldsymbol{\chi} + \boldsymbol{\varepsilon} + (1-\beta)\boldsymbol{\xi} \tag{16}$$

with $\operatorname{Var}[\boldsymbol{\varepsilon}] = \boldsymbol{\Sigma}$, and $\operatorname{Var}[\boldsymbol{\xi}] = \boldsymbol{\Xi}$. The goal is to obtain consistent estimators for β using instrumental variables or the Generalized Method of Moments (GMM). Once we have found consistent estimators, we can obtain fitted values for $\boldsymbol{k}$ and $\boldsymbol{v}$, which I will treat as the true values of k_{it} and v_{it}. I now discuss my identification strategy, based on time-series restrictions.

3.2 Time-Series Identification

We seek to obtain consistent estimators for β using restrictions on the unobservable variables k and ξ. Since my application is a panel data setting, we can tap into the large literature on internal instruments in dynamic settings (e.g., Anderson and Hsiao 1982; Arellano and Bond 1991; Blundell and Bond 1998). Specifically, I use the framework laid out in Griliches and Hausman (1986), who develop a class of estimators based on restrictions on the time-series properties of k and ξ.

To build intuition, recall that the regression of $\boldsymbol{y}$ on $\boldsymbol{x}$ is biased by the presence of the fixed effect $\boldsymbol{\gamma}$ and the endogeneity of $\boldsymbol{x}$. Assume that ξ_{it} is stationary and uncorrelated over time. Consider two common estimators developed to remove the fixed effect $\boldsymbol{\gamma}$, the within or fixed-effects estimator and the first-difference estimator:

$$\widehat{\beta}_{FE} = (\ddot{\boldsymbol{x}}'\ddot{\boldsymbol{x}})^{-1}\ddot{\boldsymbol{x}}'\ddot{\boldsymbol{y}},$$
$$\widehat{\beta}_{FD} = (\Delta\boldsymbol{x}'\Delta\boldsymbol{x})^{-1}\Delta\boldsymbol{x}'\Delta\boldsymbol{y},$$

5. Specifically, measurement error in v that is uncorrelated with k would just increase the standard errors but not affect point estimates. Hence, we can think of ψ as the component of measurement error that is common to v and k. Of course, k may have additional measurement error which would cause further attenuation bias; but as long as the time-series restrictions on ξ also apply to this k-specific measurement error, the instrumental variable strategy developed below will also eliminate this measurement error.

where $\ddot{x}_{it} := x_{it} - T^{-1}\sum_t x_{it}$, $\Delta x_{it} := x_{it} - x_{i,t-1}$ and so on. Both estimators address the fixed effect $\boldsymbol{\gamma}$, but since the regressor is endogenous with the error term, both are biased. It is not difficult to work out that as $N \longrightarrow \infty$,

$$\operatorname{plim}\widehat{\beta}_{FD} = \beta\left(1 - \frac{2\sigma_{\xi}^2}{\sigma_{\Delta x}^2}\right),$$

$$\operatorname{plim}\widehat{\beta}_{FE} = \beta\left(1 - \frac{T-1}{T}\frac{\sigma_{\xi}^2}{\sigma_{\ddot{x}}^2}\right).$$

The insight of Griliches and Hausman (1986) is to recognize that this is a system of two equations in two unknowns, β and σ_{ξ}^2; all other variances and objects are known to the econometrician once she's run the two regressions. Hence, we can solve this system to obtain a consistent estimator for β:

$$\beta = \frac{\frac{2\widehat{\beta}_{FE}}{\sigma_{\Delta x}^2} - \frac{(T-1)\widehat{\beta}_{FD}}{T\sigma_{\ddot{x}}^2}}{\frac{2}{\sigma_{\Delta x}^2} - \frac{T-1}{T}\sigma_{\ddot{x}}^2} \tag{17}$$

$$\sigma_{\xi}^2 = \frac{\left(\beta - \widehat{\beta}_{FD}\right)\sigma_{\Delta x}^2}{2\beta}. \tag{18}$$

This solution is a special case of a general estimation strategy, where the deviations between various panel-data estimators are used as instrumental variables. Formally, we seek instrumental variables estimators of the form $\widehat{\beta}_{IV} = (\boldsymbol{z}'\boldsymbol{x})^{-1}\boldsymbol{z}'\boldsymbol{y}$, where $\boldsymbol{z} = (\boldsymbol{I}_N \otimes \boldsymbol{P})\boldsymbol{x}$ for some $T \times T$ matrix $\boldsymbol{P}$. I write the following conditions for $\boldsymbol{z}$ to be a valid instrument:

Assumption 2 (Time-Series)**.**

$$\boldsymbol{\iota}'\boldsymbol{P} = 0, \tag{19}$$

$$\operatorname{plim}_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\boldsymbol{x}_i'\boldsymbol{P}\boldsymbol{u}_i = 0, \tag{20}$$

$$\operatorname{plim}_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\boldsymbol{x}_i'\boldsymbol{P}\boldsymbol{x}_i \neq 0. \tag{21}$$

The first condition ensures that the fixed effect $\boldsymbol{\gamma}$ is eliminated. The other two conditions are the usual exogeneity and relevance conditions. The instrument matrix $\boldsymbol{P}$ plays a crucial role. In the example above, we can find $\boldsymbol{P}$ by noting that equation (17) is the solution to an equation of the form for which we can rewrite the numerator as

$$\boldsymbol{x}'\boldsymbol{R}_{\mathbf{FD}}'\boldsymbol{R}_{\mathbf{FD}}\boldsymbol{y} - \boldsymbol{x}'\boldsymbol{R}_{\mathbf{FE}}'\boldsymbol{R}_{\mathbf{FE}}\boldsymbol{y} = \boldsymbol{x}'\left[\boldsymbol{R}_{\mathbf{FD}}'\boldsymbol{R}_{\mathbf{FD}} - \boldsymbol{R}_{\mathbf{FE}}'\boldsymbol{R}_{\mathbf{FE}}\right]\boldsymbol{y} =: \boldsymbol{x}'\boldsymbol{P}\boldsymbol{y} =: \boldsymbol{z}'\boldsymbol{y}.$$

In the equation above, $\boldsymbol{R}_{\mathbf{FD}}$ is a differencing matrix (a bi-diagonal matrix with -1 on the diagonal and $+1$ on the superdiagonal), and $\boldsymbol{R}_{\mathbf{FE}}$ is the within-transforming matrix, $\boldsymbol{R}_{\mathbf{FD}} = \boldsymbol{I} - \boldsymbol{J}$, where $\boldsymbol{I}$ is the identity matrix and $\boldsymbol{J}$ is $1/T$ times a matrix of all ones. In other words, the matrices $\boldsymbol{R}$ appropriately transform $\boldsymbol{x}$ to eliminate the fixed effect, and the instrument matrix $\boldsymbol{P}$ gathers the quadratic deviations of these different transformations.

I use three instruments based on this setup. My first instrument, $\boldsymbol{z}_1$, is based on the deviations between the first-difference and the within-transformation, as in the examples above. For overidentification, I also construct an instrument $\boldsymbol{z}_2$, based on the deviations between the second and first difference. The second difference is given by

$$\Delta_2 x_{it} := x_{it} - x_{i,t-2}.$$

Finally, I construct an instrument $\boldsymbol{z_3}$, based on the difference between the within-estimator and the second-difference estimator. I derive in Appendix A that the vectors of instruments $\boldsymbol{z}$ take the form

$$\boldsymbol{z_1} = \begin{pmatrix} \bar{x} - x_2 \\ \bar{x} - x_1 + x_2 - x_3 \\ \bar{x} - x_2 + x_3 - x_4 \\ \vdots \\ \bar{x} - x_{T-1} \end{pmatrix}, \boldsymbol{z_2} = \begin{pmatrix} x_1 - 2x_2 + x_3 \\ -2x_1 + 4x_2 - 3x_2 + x_4 \\ x_1 - 3x_2 + 4x_3 - 3x_4 + x_5 \\ \vdots \\ x_{T-3} - 3x_{T-2} + 4x_{T-1} - 2x_T \\ x_{T-2} - 2x_{T-1} + x_T \end{pmatrix}, \boldsymbol{z_3} = \begin{pmatrix} \bar{x} + x_1 - 3x_2 + x_3 \\ \bar{x} - 3x_1 + 5x_2 - 4x_3 \\ \bar{x} - 3x_2 + 5x_3 - 4x_4 \\ \vdots \\ \bar{x} - 3x_{T-1} + x_T \end{pmatrix}. \tag{22}$$

where $\bar{x} := T^{-1} \sum_t x_{it}$ is the time-series average of x for firm i. This setup is easily extended to unbalanced panels by making the period length observation-specific, i.e., T_i instead of T.

In principle, I could use more instruments to improve efficiency. The optimal number depends on the nature of the restrictions I impose on $\boldsymbol{k}$ and $\boldsymbol{\xi}$. The previous instruments were derived under the assumption that $\boldsymbol{\xi}$ is stationary and uncorrelated over time. This may be a strong assumption, but given the overidentification, it is testable. Griliches and Hausman (1986) derive the optimal number of instruments as the largest number which all contain independent information about V and K. The maximal number of instruments is T^2. The optimal number of instruments is therefore T^2 minus the number of unique linear restrictions provided by equations (19) and (20) (since (21) is typically not binding). Requirement (19) imposes T restrictions. The optimal number of instruments further depends on the time-series properties of the measurement errors. To see this, rewrite (20) as

$$\operatorname{plim}_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} (\boldsymbol{k}_i + \boldsymbol{\xi}_i)' \boldsymbol{P} (\boldsymbol{\varepsilon}_i - \boldsymbol{\xi}_i (1+\beta)) = 0 \Rightarrow$$
$$\operatorname{plim}_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{\tau=1}^{T} \xi_{it} \xi_{i\tau} P_{i\tau} = 0.$$

The implied restrictions on $\boldsymbol{P}$ depend on Ξ, the covariance matrix of $\boldsymbol{\xi}_i$. If ξ_{it} are not stationary nor serially correlated, the exogeneity condition requires that $p_{t\tau}$ equal zero whenever $t = \tau$. This imposes T additional restrictions on $\boldsymbol{P}$. If ξ_{it} are stationary, then only $\mathrm{tr}(\boldsymbol{P})$ must equal 0. This is one additional restriction. If the measurement errors follow a MA(m) process, this imposes restrictions of either $m+1$ (if $m < T-2$) or $T-1$ (if $m \geq T-2$).

In my main application, I will assume that ξ_{it} is stationary and serially uncorrelated; under these conditions, the instrument vectors $\boldsymbol{z_1}, \boldsymbol{z_2}$ and $\boldsymbol{z_3}$ found above are consistent. In extensions, I relax the assumptions on the covariance structure of ξ by allowing for MA(1) dependence and hence autocorrelation. I maintain stationarity throughout.

4 Data

I use administrative data from Statistics Netherlands. I merge several large datasets. I first describe the firm data I use, and then the household data.

4.1 Discount Rates

To construct my estimate of market value, y_{it}, I need discount rates. Since firm equity is a risky asset, the discount rate must be risky. Constructing a discount rate for non-listed assets is challenging, however, since both the costs for equity and for debt are unobserved (Damodaran 2012). For listed assets, one could use one of the various asset pricing models (e.g., Fama and French 1993). Bach, Calvet, and Sodini (2020), for instance, obtain expected returns in private equity by estimating such an asset pricing model on listed stocks, regressing the obtained risk factors on listed firm characteristics, and then assuming that these coefficients are the same for unlisted businesses. These are strong assumptions that strongly depend on the chosen asset-pricing model.

I use three different measures of the discount rate, each following different strands of the literature. The first measure is based on the conceptually correct discount rate, the weighted-average cost of capital:

$$\rho_t^{\text{wacc}} = \omega_t \times (1-\tau) \times r_t^d + (1-\omega_t) \times r_t^e, \tag{23}$$

where ω is the leverage ratio (i.e, the share of debt in total firm liabilities), τ is the firm's tax rate, and r^d and r^e are the costs of debt and equity. Computing the WACC for private firms is again challenging. I opt for an indirect approach, using data by Gormsen and Huber (2023, 2024), who calculate the WACC for a large sample of firms based on rich asset-pricing models. The firms in their sample are listed firms who make conference calls, in which they also mention the firms' self-reported discount rates. Gormsen and Huber find that these internal discount rates persistently differ from the WACC. For my first two measures, I take the Dutch firms from their sample, and simply use the average of their data per year. This results in a series ρ^{wacc}, based on the weighted-average cost of capital, and a series ρ^{gh}, based on firms' self-reported discount rates.

As a final measure, I follow a large stream of literature, including Barkai (2020), and estimate ρ^{b} as the ten-year Dutch government bond yield plus a constant risk premium of 5%. Since bond yields have declined strongly, this measure of ρ declines as well, which will blow up valuations.

Figure 1: Discount Rates

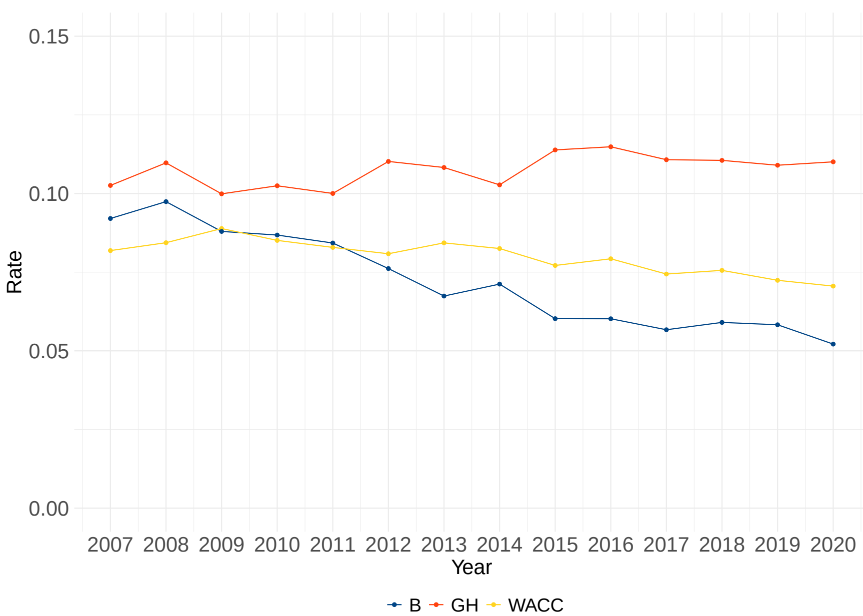

Figure 1 plots the different measures. The three discount rates comove, but also display important differences. The series for ρ^{b} declines strongly following the financial crisis in 2008, whereas the other two series remain much more stable. The self-reported discount rates, ρ^{gh} are persistently the highest and stable around 11%. The weighted-average cost of capital, which is conceptually the closest to the true discount rate, is in between these two series for most of the period.

Obviously, all three discount rate measures are highly imperfect; this is why my measurement-error framework is important. Under the identifying assumptions of the model, my GMM estimates will retain only those parts of these series that are structurally related to firms' capital stock. Ideally, the instruments should lead to similar point estimates across the series, since that gives confidence that our initial choice of discount rate does not drive the estimation results.

4.2 Firm Data

My main data source is the *Bedrijfsgegevensstatistiek* (BG), which covers the universe of Dutch corporate income tax returns 2007–2019, where dates are December 31st of a given year. Hence, all my results apply to incorporated firms.[6] The Netherlands levies a progressive corporate income tax with two brackets. For most of my sample period, the first bracket applied to fiscal profits up to € 200,000 (with a rate of 20%), and the second bracket applied to all profits in excess of € 200,000 (taxed at 25%). From the corporate income tax returns, I obtain the full balance sheet and profit and loss statement of each firm. In addition, I observe firm characteristics such as industry, corporate form, and firm age, although coverage varies for these background variables.

The predominant corporate form in my data is the non-listed limited-liability company (*besloten vennootschap*, BV). This is a widely used form that covers a wide variety of firms. In particular, it is not uncommon to observe nests of firms, where firm A owns firm B which owns firms C and so on. The BG data shows balance sheet and income statement variables both at the consolidated and unconsolidated level. Since my ultimate interest is to link firms to owners, I use consolidated data throughout.

For firm's capital $\boldsymbol{x}$, I use the sum of the firm's fixed capital stock and intangible capital. This choice of variables is done for substantive reasons. The neoclassical investment model emphasizes the role of productive capital in firms' investment decisions. Hence, while balance sheet items such as property contribute to a firm's value, they are not part of the firm's productive capital stock.

In Dutch accounting rules, intangible assets include goodwill, development (not research), software, and so on. There are two criteria to include intangible assets on the balance sheet. First, it must be likely that the asset generates future economic benefits. Second, the costs of the asset can be reliably assessed. Some costs do not meet these criteria and are instead expensed on the income statement, such as costs for training, advertisement, startup costs for a new product or activity. Research is expensed, not an asset; development can be included as an asset if the to be developed intangible asset is (a) technically feasible; (b) the intention exists to use or sell the asset; (c) the capacity to produce/use the asset must be there; (d) it must be likely to generate future economic benefits; (e) the capacity is there to complete the asset; (f) the costs can be reliably assessed.

To construct my measure of firms' market value, $\boldsymbol{y}$, I combine the discount rates from the previous section with a moving average of firm profits. I use the three-year moving average as a smoothed forecast of future profits.[7] Figure 2 plots the aggregate trends of this measure, $\overline{\Pi}_t$. The series is mostly stable, hovering between € 12.5 billion and € 17.5 billion for most of my sample period. In the first year, averaged profits were higher, with the 2008 value peaking at around € 22.5 billion.

My estimation methods rely on existing values for $\boldsymbol{x}$; hence, I drop all observations where this variable is missing. This results in a dataset with about 8 million firm-year observations. Note that not all these observations can be linked to Dutch households, to which I turn next.

4.3 Household Data

I use administrative data on Dutch households, which are available from 2006–2020, measured on January 1st.[8] This dataset covers households' total wealth and its composition. We can distinguish between deposits, financial assets (which are stocks,

6. Note that partnerships and other unincorporated business forms are less prevalent in the Netherlands among the top of the distribution than in other countries, such as the United States (Kopczuk and Zwick 2020). This is because unincorporated businesses are subject to the personal income tax, whereas incorporated firms have distinct tax advantages, as I detail in the next section.

7. In Crouzet and Eberly (2023)'s model, Π_t refers to a firm's operating surplus, i.e., total revenue minus variable costs. This variable is less well observed in my data than profits, and manual reconstructions result in noisier and more volatile series than simply using profits. I return to this issue when estimating the model's structural parameters in Section 6.

8. Hence, when linking firm and household data, I assign firm data from December 31st in year t to year $t + 1$.

Figure 2: Aggregate Trends in $\overline{\Pi}_t$

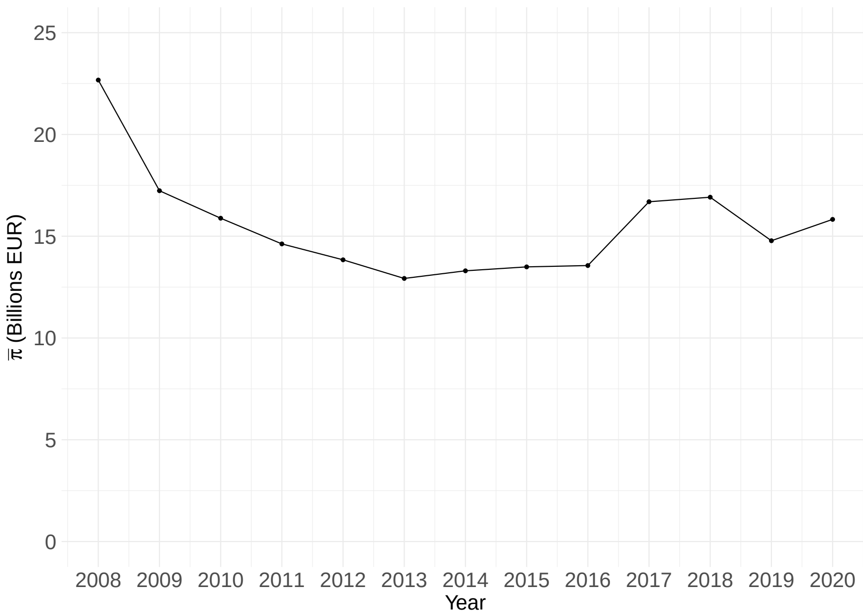

bonds, and other securities), owner-occupied housing, other real estate, non-incorporated business wealth, incorporated business wealth, other assets, mortgages, student loans, and other liabilities. Since capital-funded pensions are not taxed, they are excluded from all wealth measures.[9]

The Netherlands levies different types of taxes on different types of assets under its income tax system, which is organized in three 'boxes'. The imputed rent from owner-occupied housing is taxed together with labor earnings (including earnings of the self-employed) in Box 1. This imputed rent increases in the cadastral value of the property, making this Box nominally progressive in capital income; however, since the interest on mortgage debt is deductible, net taxable income from owner-occupied housing is typically negative.[10] Box 2 taxes dividends and realized capital gains from corporations in which a household has more than 5% of the shares (*aanmerkelijk belang*, lit. "significant stake") This tax is progressive and is paid on top of the corporate income tax described in the previous section.

All other assets (except pension claims) are taxed in Box 3, which taxes a presumptive or fictitious rate of return.[11] For most of my sample period, Box 3 charged a flat tax of 30% on a presumptive return of 4% (implying an effective wealth tax of 1.2%), above some threshold. From 2017, progressive rates were introduced in Box 3, based on a presumed portfolio composition at different wealth levels, where the presumptive return in each bracket was the weighted average of the historical return of these assets (deposits, stocks, bonds, etc.).[12]

9. It is debatable whether this omission is justified. From a permanent-income perspective, funded pension claims enter the net present value of future consumption and hence should be included in wealth. However, the Dutch pension system is different from comparable systems in countries such as the United States in that individuals cannot consume, trade, or otherwise claim their pension assets before retirement; in other words, households hold no property rights over their pension claims. This makes the Dutch capital-funded system more akin to Social Security, which are typically not included in household wealth. In incomplete financial markets, there is a meaningful distinction between present-value-based and property-rights-based definitions of wealth. See Martínez-Toledano, Sodano, and Toussaint (2023) for further discussion of these points and their implications.

10. The Netherlands has uniquely large mortgage debt, in the range of 100% of national income in the period I cover (Toussaint, de Vicq, Moatsos, and van der Valk 2022). This has been driven by government stimulus of owner-occupied housing, as exemplified by interest-only mortgages (Bernstein and Koudijs 2024).

11. Since legally the tax is on a return to assets, it is called a capital income tax. Since the rate of return is not the realized return but a presumptive return, Box 3 is a *de facto* wealth tax (Jacobs 2015).

12. This differentiation was challenged in court and declared illegal by the Supreme Court in 2022, who held that it discriminated against individuals who did not hold this exact portfolio composition. As a result of this ruling, the current Box 3 system based on presumptive taxation has been declared illegal and the government is obligated to move towards a system of taxing realized returns.

This wide variation in tax rates according to asset class leads to wide variation in effective tax rates on households' assets, and therefore also provides fiscal incentives to move assets to low-tax Boxes. In particular, Box 2 provides incentives to defer taxation by deferring dividend payouts (although the corporate income tax continues to apply). Firm-owners can borrow against the value of their firm and use these loans to finance their consumption; since these loans depress the value of their assets, this might reduce additionally reduce their tax burden in Box 3 (IBO Vermogensverdeling 2022). Indeed, in their analysis of Dutch tax progressivity, Bruil et al. (2022) find that the Dutch overall tax system is regressive, which is largely due to the fiscally advantageous role of Box 2.

Since wealth is typically owned jointly by households, all data are on the household level. We link this dataset to firm data using the Shareholder Registry (*Aandeelhoudersregister* (AR)), see Appendix B. This dataset provides information on the ownership structure of each incorporated firm at quarterly frequency. Linking the datasets results in about 7 million firm-owner-year observations overall. I merge this linked set with the full dataset on the Dutch wealth distribution, ending up with more than 120 million observations at the household level, about 7 million of which therefore are firm-owner observations.

4.4 Summary Statistics

Table 1 provides summary statistics for all firms which are matched to Dutch household owners. It provides selected items from the balance sheets and income statements, as well as ownership concentration.

Table 1: Summary Statistics, Firms & Firm Owners

Item	Min	P25	Median	Mean	P75	Max
Panel A: Balance Sheet						
Intangible Capital	0	0	0	0.031	0	6,217
Property	-0.368	0	0	0.228	0	1,824.529
Fixed Capital	-0.685	0	0.001	0.321	0.048	3,394.027
Shares	-1,564	0	0	0.834	0.018	44,670
Inventory	-296.648	0	0	0.750	0	3,831.089
Claims	-8.58	0.008	0.043	0.475	0.159	235,900
Securities	0	0	0	0.233	0	9,654
Total Assets	-1,552	0.102	0.310	2.509	0.902	243,000
Long-Run Liabilities	-0.179	0	0	0.294	0	21,970
Short-Run Liabilities	-2,475	0.005	0.025	0.557	0.109	16,680
Net Worth	-853.8	0.008	0.089	1.367	0.437	45,650
Panel B: Income Statement						
Net Revenue	-7.265	0	0.045	0.959	0.198	227,800
Personnel Costs	-16.559	0	0.03	0.174	0.096	508.922
Depreciation	-9.651	0	0	0.008	0	43.216
Total Costs	-91.1	0.004	0.057	0.841	0.174	78,530
Net Financial Result	-250.342	-0.003	0	-0.002	0.005	2,015.167
Shareholdings Result	-7,173	0	0	0.098	0	5,265
Fiscal Profit	-884.602	-0.004	0.003	0.065	0.035	4,341.089
Panel C: Ownership						
Shares (%)	0.128	50	100	76.563	100	100

Notes: All observations pooled (6,319,593 observations); values in millions of nominal EUR. Not all items on the balance sheet and profit & loss statement are summarized above.

It is clear that firm values are widely dispersed. Some firms are worth in excess of € 45 billion, while the median firm is worth only € 89,000. Similarly, there is wide variation in the structure of firms' balance sheets. Half of all firms have negative or negligible values of capital stocks, and similarly for most other components. The income statement also shows wide variation. Personnel costs and depreciation are less dispersed than financial variables such as the firm's capital income flows. Panel C reveals that most firms are fully owned by a single owner. Given the large values of firm wealth, this suggests that firms play an outsize role in wealth concentration.

This suggestion is confirmed by Table 2, which dissects the wealth distribution for the year 2019. Note that the summary statistics in this Table are all at the household level.

Table 2: Wealth Distribution, 2019

					Composition (Share of Assets)				
Bracket	# Households	Threshold	Mean	Share	Deposits	Financial	Real Estate	Business	Debt
Population	8,045,662		212,638	1.000	0.125	0.071	0.622	0.182	−0.334
Bottom 50%	4,022,825		−1,907	−0.004	0.152	0.012	0.830	0.006	−1.033
Middle 40%	3,218,268	48,838	197,988	0.372	0.134	0.023	0.815	0.029	−0.404
Top 10%	804,569	466,418	1,343,958	0.632	0.113	0.124	0.420	0.344	−0.145
Top 1%	80,457	2,197,062	5,935,055	0.279	0.056	0.148	0.222	0.574	−0.123
Top 0.1%	8,046	9,749,227	24,404,472	0.115	0.033	0.139	0.127	0.700	−0.099
Top 0.01%	805	40,537,908	93,433,157	0.044	0.023	0.116	0.067	0.794	−0.064

Notes: Values for the threshold and mean in nominal EUR. Financial assets include stocks, bonds, and other securities. Real estate is the sum of owner-occupied housing and other real estate. Business is the sum of incorporated and non-incorporated business wealth. Debt is the sum of mortgages, student loans and other liabilities.

We first note that wealth is highly concentrated. The top 1% of households owns almost a third of all wealth. The top 0.01% – 805 households – own 5%. Another important observation is that the composition of wealth varies dramatically over the wealth distribution. The bottom 90% have most of their wealth in housing and deposits, differing only in the relative value of their house vis-a-vis their mortgage. Only in the top 10% do financial assets start to play a meaningful role. In fact, most of the top 10% looks similar in wealth composition to the bottom 90%, and only from the top 1% or so upwards do we see financial and business wealth become prominent. Of the two, the latter is by far the most important. The top 0.01% hold almost 80% of their assets in private business wealth. This confirms that the correct valuation of these assets is of key importance for studying wealth concentration at the top.

5 Estimation Results

In this section, I report the results from my econometric procedure. I begin with simple OLS estimates to give a benchmark. Then, I move to the main results of the paper. I finish the section by exploring some extensions.

5.1 OLS

We begin our investigation by simple OLS regressions. Recall that I use three distinct discount rates: ρ^{wacc}, ρ^{gh}, and ρ^{b}. I use these discount rates to capitalize smoothed profits, $\overline{\pi}_{it}$, to create three distinct raw series of firm value, y_{it}. In Table 3, I regress these series on the firms' capital stock. I report Driscoll-Kraay standard errors, which are robust to arbitrary serial correlation and heteroskedasticity.

I first investigate a simple regression, in columns (1), (4), and (7). Then, in succeeding columns, I add fixed effects for year and firm sector (at the two-digit level). Across specifications, coefficients remain stable and are always highly significant. Adding fixed effects makes no difference. To some extent, this is to be expected: our measures of y_{it} are based on the same

Table 3: OLS Estimates

	ρ^{wacc}			ρ^{gh}			ρ^{b}		
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
β	4.23***	4.23***	4.24***	3.37***	3.37***	3.37***	4.54***	4.55***	4.55***
	(0.664)	(0.665)	(0.665)	(0.576)	(0.576)	(0.577)	(0.533)	(0.533)	(0.533)
Fixed-effects									
Year		✓	✓		✓	✓		✓	✓
Industry			✓			✓			✓
Fit statistics									
N	6,713,564	6,713,564	6,713,564	6,713,564	6,713,564	6,713,564	6,713,564	6,713,564	6,713,564
R^2	0.335	0.335	0.335	0.335	0.335	0.335	0.340	0.340	0.340
Within R^2		0.335	0.335		0.335	0.335		0.340	0.340

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*
Notes: Driscoll-Kraay standard-errors in parentheses.

discount rate for all firms within a year; hence, all variation must be within-year variation. Nevertheless, it is surprising that adding industry fixed effects does not alter any of the results. Based on these specifications, we would expect a coefficient of β in the region of 3.5–4.5. However, the measurement error concerns noted earlier imply that the true value could well be outside this interval too. Next, I investigate whether using instruments in a GMM procedure improves identification.

5.2 Time-Series Identification

Next, we look at the time-series results. Table 4 shows the manual identification of β using equations (17) and (18). Column (1) shows the regression of the within estimator, and (2) that of the first-difference estimator. Note that the identification of Griliches and Hausman (1986) relies on both estimators having the same numbers of observations. Therefore, whenever a lagged value is not available, I set that value to 0 instead of missing. This is consistent with the econometric framework developed in Appendix A.2.

Both values are positive, but again insignificant. Note that $\widehat{\beta}_{FE} > \widehat{\beta}_{FD}$, which is to be expected; Griliches and Hausman (1986) argue that this indicates a panel structure with a declining autocorrelogram in the variables. In the bottom panel, we identify $\beta \in (2.1, 4.4)$, using the estimated variances of the regressions and point estimates from the three measures. This is a wide interval, but again, given the insignificance of the point estimates, it is unclear where in this interval the true value lies. The manual calculations underlying this Table do not allow to weight the data, which results in a loss of efficiency.

I now systematically employ the instrumental variables strategy based on the time-series restrictions. I use two-step GMM (Hansen 1982), which is asymptotically consistent. Moreover, since the weighting matrix converges to the optimal weighting matrix, the results should be more efficient than either the OLS or manual calculations reported above. I use heteroskedasticity and autocorrelation-robust standard errors, based on Andrews (1991).[13] The results are in Table 5.

Table 5 consists of three panels, one for each measure for y_{it}. Each panel has seven columns, which each report estimation results based on all possible combinations of my three instrumental variables. I report robust first-stage F-statistics based on Montiel Olea and Pflueger (2013), which are robust to serial correlation, clustering, and heteroskedasticity. The endogenous right-hand side variable is x_{it}, which is identical across all 21 specifications; hence, it should be no surprise that the first-stage F-statistic does not vary between panels in the same column.

Start with the just-identified columns (1)–(3). Column (1) is insignificant across panels. Columns (2) and (3) are significant

13. I use the Bartlett kernel, with a bandwidth of three.

Table 4: Within vs. First-Difference Transformation

Discount Rate:	r^{wacc}		r^{gh}		r^{b}	
Dependent Variable:	$\ddot{y}_{it}$	Δy_{it}	$\ddot{y}_{it}$	Δy_{it}	$\ddot{y}_{it}$	Δy_{it}
Model:	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{\beta}$	2.92	1.81	2.06	1.43	4.02	1.84
	(2.46)	(1.59)	(1.78)	(1.28)	(3.03)	(1.53)
Fit statistics						
Observations	6,713,564	6,713,564	6,713,564	6,713,564	6,713,564	6,713,564
R^2	0.788	0.062	0.799	0.069	0.773	0.052
Within R^2	0.093		0.078		0.138	
Separate Parameter Estimates:						
$\sigma^2_{\ddot{x}}$	2.23×10^{15}		1.35×10^{15}		2.72×10^{15}	
$\sigma^2_{\Delta x}$		6.65×10^{14}		3.70×10^{14}		8.33×10^{14}
Joint Parameter Estimates:						
β	3.09		2.15		4.38	
σ^2_{ξ}	1.38×10^{14}		6.21×10^{13}		2.41×10^{14}	

Driscoll-Kraay standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

and the coefficients are sizable.The robust first-stage F-statistics are quite sizable for those columns, although strictly below the thresholds where the asymptotic bias of the GMM estimator is no more than 10% of the OLS bias (which is the criterion used by Montiel Olea and Pflueger 2013). Column (2) is significant at the 5% level in all panels, whereas column (3) is significant at the 0.1% level. Moreover, coefficients in column (3) are generally larger than in column (2), with the exception of Panel C.

Next, we move to the specifications with more instruments than endogenous variables, which permit overidentification tests. Columns (4)–(7) show J-tests which are uniformly insignificant. In columns (4) and (6), these test statistics are actually almost equal to 0. The statistics in column (7) are larger in magnitude, but given the extra degree of freedom these statistics are actually very insignificant. These results indicate that the identifying assumptions underpinning my GMM procedure are easily upheld. Coefficients in these columns tend to cluster closer together within each panel, although there is quite some variation across panels. The coefficients for ρ^{wacc} are in the 3.9–4.5 range, while those for ρ^{gh} cluster between 2.7–3.6 and those for ρ^{b} range between 4.6 and 5.0. Coefficients in column (4) are significant at the 5% level; those in (5)–(7) are significant at the 0.1% level.

The first-stage F-statistic in columns (4)–(5) is reduced compared to columns (3). This indicates that these combinations of instruments does not add additional information that can be used to significantly identify the first-stage coefficient. Column (6), on the other hand, has higher F-statistics than (3), and column (7) is comparable in magnitude. These coefficients also do not meet the Montiel Olea and Pflueger (2013) threshold above which an instrumental-variables estimate is asymptotically less biased than the comparable OLS estimate; although it should be noted that the standard Kleibergen-Paap F-statistic is much higher for column (7), on the order of 43. Thus, formally my coefficients are estimated with weak instruments (Andrews, Stock, and Sun 2019); nevertheless, they appear reasonably informative. Moreover, I agree with Andrews, Stock, and Sun (2019) that statistical weakness does not equal economic weakness; my instruments are derived under a structural model that is quite general (Crouzet and Eberly 2021). Hence, on substantive grounds, my instruments should be valid.

Column (6) combines z_1 and z_3. Interestingly, the coefficients for ρ^{wacc} and ρ^{b} in Panels A and C are now almost identical around 4.6, where they diverged quite substantially in specifications (1)–(6). The coefficient in Panel B remains lower at 3.6.

Table 5: GMM Estimation, Time-Series Identification

Dependent Variable:	y_{it}						
Instruments	z_1	z_2	z_3	z_1, z_2	z_2, z_3	z_1, z_3	z_1, z_2, z_3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: ρ^{wacc}							
β	4.610	4.089*	4.543***	3.882*	4.530***	4.529***	4.347***
	(3.077)	(2.001)	(1.238)	(1.865)	(1.237)	(1.141)	(1.031)
Observations	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563
J-test Statistic				0.093	0.078	0.001	0.159
J-test p-value				0.761	0.780	0.977	0.923
First Stage F-statistic	1.934	6.205	7.785	2.932	6.507	8.658	7.066
Panel B: ρ^{gh}							
β	3.242	2.909*	3.580***	2.782*	3.507***	3.640***	3.423***
	(2.221)	(1.433)	(0.977)	(1.337)	(0.962)	(0.929)	(0.828)
Observations	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563
J-test Statistic				0.070	0.303	0.038	0.338
J-test p-value				0.792	0.582	0.845	0.845
First Stage F-statistic	1.934	6.205	7.785	2.932	6.507	8.658	7.066
Panel C: ρ^{b}							
β	6.347	5.289*	5.052***	4.845*	5.031***	4.743***	4.615***
	(3.940)	(2.590)	(1.393)	(2.383)	(1.380)	(1.177)	(1.059)
Observations	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563
J-test Statistic				0.255	0.014	0.176	0.263
J-test p-value				0.614	0.907	0.675	0.877
First Stage F-statistic	1.934	6.205	7.785	2.932	6.507	8.658	7.066

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Two-step GMM with heteroskedasticity and autocorrelation-robust standard errors in parentheses. Robust first-stage F-statistic by Montiel Olea and Pflueger (2013). Each panel shows GMM results with as independent variable the $\boldsymbol{y}$ computed using the respective discount rate.

Column (7) combines all three instruments. The resulting coefficients are again very similar in panels A and C, and consistent with the coefficients in column (6). The coefficient in Panel B remains lower, but is very close to that in column (6) as well.

I conclude that my GMM procedure holds up. The framework results in coefficients which are reasonable in magnitude, highly significant, with reasonably strong first stages and insignificant overidentification tests. The fact that coefficients for ρ^{wacc} and ρ^{b} converge, despite being based on very different discount rate trends and despite having very different OLS coefficients (see Table 3), gives confidence that these coefficients are close to the true value. On the other hand, from a data-driven perspective there is no reason to reject the coefficients in Panel B. Obviously, the lower coefficients in Panel B will result in lower values for private business wealth and hence inequality.

In the end, I use two coefficients. The first, dubbed $\beta^{+} \approx 4.56$, is the average of columns (6) and (7), Panels A and C. The second, $\beta^{-} \approx 3.53$, is the average of columns (6) and (7) in Panel B. With this high and low coefficient, I will estimate top wealth shares and heterogeneous returns to wealth. In Section 6, I investigate whether the magnitudes of these coefficients make economic sense at the aggregate level.

5.3 Extensions

So far, the time-series identification rested on the identifying assumption that the measurement errors ξ_{it} were stationary and serially uncorrelated. Griliches and Hausman (1986) show that it is possible to relax these assumptions and still obtain valid instruments. I now assume that the measurement errors are stationary but follow an MA(1) process. The strategy is, as before, to find an estimating framework that gets rid of the fixed effect $\boldsymbol{\gamma}$, and finding valid instruments to that effect. Following Griliches and Hausman (1986), we end up with the following system of three equations:

$$\Delta y_{it} = \Delta x_{it} + \Delta u_{it} \tag{24}$$

$$\Delta_2 y_{it} = \Delta_2 x_{it} + \Delta_2 u_{it} \tag{25}$$

$$\Delta_3 y_{it} = \Delta_3 x_{it} + \Delta_3 u_{it} \tag{26}$$

where $\Delta_\tau y_{it} := y_{it} - y_{i,t-\tau}$, i.e., the long difference of order τ. Under the assumption that ξ is stationary and follows an MA(1) process, valid instruments for the equations are, respectively:

$$z_4 = \left(x_{i,t-1} + x_{it}, \quad x_{i,t+2}, \quad x_{i,t-2}\right)' \tag{27}$$

$$z_5 = \left(x_{i,t-1}, \quad x_{i,t+1}\right)' \tag{28}$$

$$z_6 = \left(x_{i,t-2}, \quad x_{i,t-1}, \quad x_{i,t+1}\right)' \tag{29}$$

We simultaneously estimate this system of equations using system-GMM. Table 6 reports the results.

The table shows a value for β that is positive and mostly uniform across the equations, but lower in magnitude than the results obtained previously. Moreover, the coefficients are all insignificant, indicating that there is insufficient within-firm variation in long differences to identify β (standard errors are clustered at the firm level). The robust first-stage F-statistics of these models are insignificant except for the third equation, which is highly significant. The joint J-test is insignificant, indicating that the over-identifying assumptions fail to be rejected.

Since the coefficients are insignificant, it is unclear whether the true coefficient is lower than those found in Table 5, since the 'consensus' coefficients from that Table are within a standard error from the coefficients reported here. Moreover, these specifications are in first and longer differences, whereas the earlier results are for variables in levels. Given the lack of unambiguous differences relative to Table 5, I stick with the coefficients from that table.

Table 6: System GMM Estimation, Time-Series Identification

	ρ^{wacc}			ρ^{gh}			ρ^{b}		
$\boldsymbol{y}$	Δy_{it}	$\Delta_2 y_{it}$	$\Delta_3 y_{it}$	Δy_{it}	$\Delta_2 y_{it}$	$\Delta_3 y_{it}$	Δy_{it}	$\Delta_2 y_{it}$	$\Delta_3 y_{it}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
β	1.937	1.895	1.583	1.509	1.488	1.236	2.206	1.996	1.653
	(1.841)	(1.175)	(0.997)	(1.411)	(0.934)	(0.802)	(2.033)	(1.183)	(0.981)
$\boldsymbol{x}$	Δx_{it}	$\Delta_2 x_{it}$	$\Delta_3 x_{it}$	Δx_{it}	$\Delta_2 x_{it}$	$\Delta_3 x_{it}$	Δx_{it}	$\Delta_2 x_{it}$	$\Delta_3 x_{it}$
z	z_4	z_5	z_6	z_4	z_5	z_6	z_4	z_5	z_6
N	6,713,561	6,713,561	6,713,561	6,713,561	6,713,561	6,713,561	6,713,561	6,713,561	6,713,561
F-test	2.621	5.205	14.862	2.621	5.205	14.862	2.621	5.205	14.862
J-test		2.613			2.766			2.201	
p-value		0.759			0.736			0.821	

6 Aggregate Implications

In this Section, I investigate the aggregate implications from my empirical results. First, I contrast my alternative measures of firm value to the recorded book values. Second, I investigate whether the data are consistent with the model laid out in Section 2. Both exercises allow me to check whether the results obtained in the previous Section make economic sense.

6.1 Aggregate Values

With the results from Section 5 in hand, we can estimate the changes in private firm value. I first present estimates at the aggregate level. To do so, I aggregate three series: the microdata (based on book values, denoted b_{it}), the raw estimated market values based on a discount rate (y_{it}), and the corrected market values based on the GMM procedure ($\widehat{y}_{it} =: v_{it}$). I then aggregate these as

$$B_t = \sum_i b_{it},$$
$$Y_t = \sum_i y_{it},$$
$$V_t = \sum_i v_{it}.$$

In Figure 3, I report the results of this aggregation. I use y^{wacc} and I use the fitted values using β^{+}. When moving to wealth shares in the next section, I more systematically compare different fitted values. I also show values for the aggregate capital stock $K_t = \sum_i k_{it}$.

The microdata, based on book values B_t, show a steep upward trend. In contrast, the raw alternative measure Y_t is relatively stable, only increasing somewhat after 2016. It is also generally lower than B_t. The relative stability is a function of both the relative stability of $\overline{\Pi}_t$ and ρ^{wacc}. Figure 3 also shows the importance of adjusting the raw Y_t using the GMM procedure. Effectively, V_t is the component of Y_t that is structurally related to the true capital stock K_t and is not attenuated by measurement error. As a result, the aggregate updated series for private business market value, V_t, is generally larger than both the book value and Y_t. At the beginning of the sample, it exceeds B_t by around €150 billion, which is on the order of 30% of GDP. It is also increasing over time, but at a slower rate than the book value, reflecting that part of the steep increase in B_t is an accounting artefact rather than a true increase in economic value. After 2016, the two series more ore less coincide. In the

Figure 3: Adjusted Values for Total Private Business Wealth

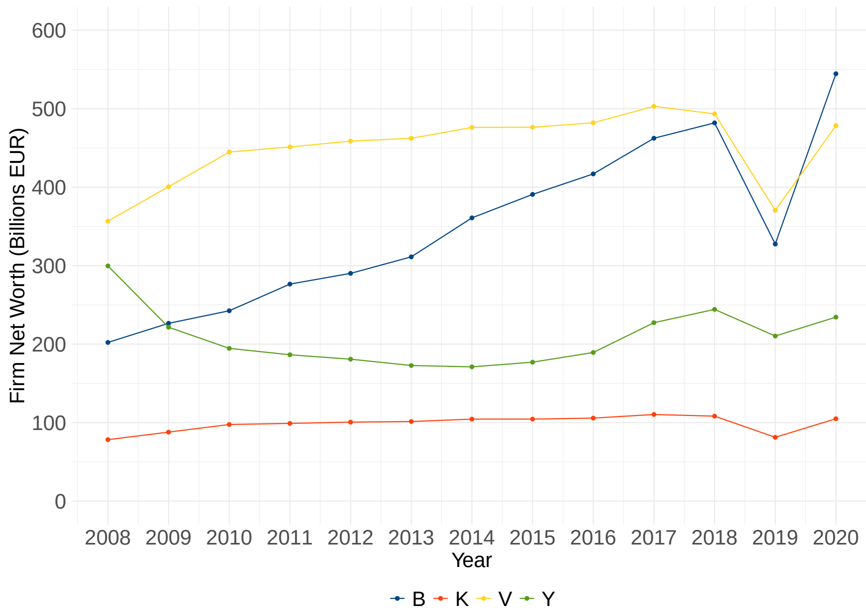

Notes: All values are aggregated across firms. B = book value (recorded private business wealth); K is the capital stock, Y is the raw alternative measure of firm value (based on $\overline{\Pi}_t$ and using the ρ^{wacc} discount rate); and V is the fitted value from Table 5 (using β^{+}).

final two years, the book value outpaces my market-value measure. This is reassuring, since it means that my adjustments do not mechanically increase market values. Instead, the updated V_t seems to be a different series than B_t (although obviously correlated), lending confidence to our results.

The steep increase in B_t is puzzling, because many of the fundamentals seem not to have changed. The reported capital stock K_t hovers around €100 billion, and Figure 2 shows that aggregate profits are also stable. Of course, profits accumulate into book value, so stable values for $\overline{\Pi}_t$ would result in a constant increase in B_t. The steepness of the increase, however, cannot solely be explained by profit accumulation. Inspection of the balance sheets suggests that a substantial part of the increase might be driven by fiscal motives. As explained in Section 4, financial assets other than private businesses are subject to a wealth tax in Box 3, based on a presumptive rate of return. As interest rates declined, this presumptive rate of return became increasingly difficult to achieve, especially for deposits and other safe assets. Reallocating toward private businesses allowed investors to escape the increasing tax burden in Box 3.

The granular data of the corporate income tax returns allows me to partially investigate this hypothesis. Figure 4 shows the time series of aggregate gross capital inflows into firms, i.e., capital reallocated towards the firm by its owners. This variable is distinct from profits and is purely a financial reallocation. I use gross inflows because I am interested in an allocation towards private business firms, which are taxed in the fiscally advantageous Box 2. Using net inflows (where I correct for outflows) makes little difference to these results.

Figure 4 shows that inflows were substantial in this period. Moreover, after 2013 inflows increased substantially. The series is quite volatile, but the post-2013 average is clearly higher than the pre-2013 average. This coincides with the increase in the gradient of the B_t series in Figure 3. The two peaks in 2014 and 2017, moreover, coincide with the steepest increases in B_t. This suggests that capital inflows are a substantial explanation for the increase in recorded book values. The post-2013 period also saw a sustained decline in global safe interest rates, as can be gleaned from the ρ^{b} series in Figure 1. Of course, safe interest rates were already declining before 2013 as well. However, bank deposit interest rates declined steeply after 2013, as can be

Figure 4: Capital Inflows, 2008–2020

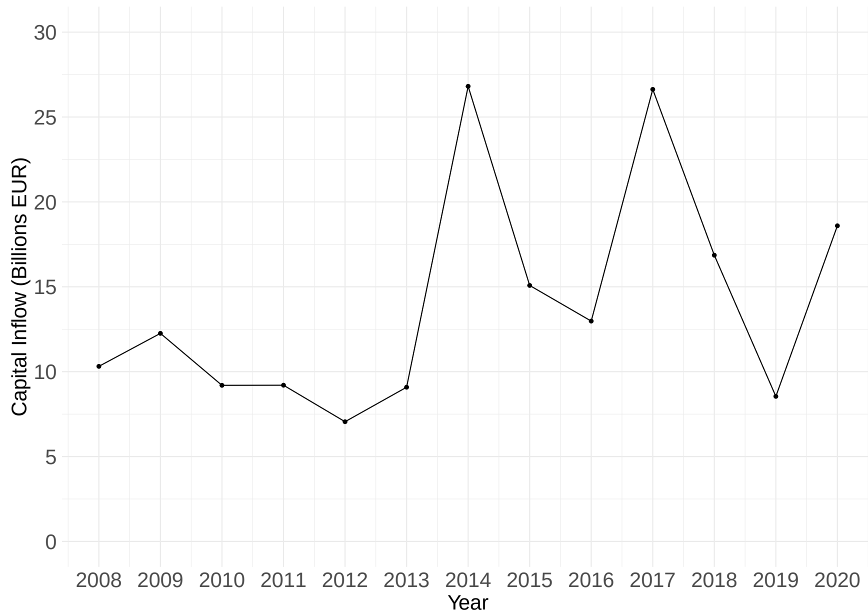

Notes: Figure shows aggregate values for gross capital inflows into firms.

seen in Figure 5. This Figure plots monthly deposit interest rates for Dutch account-holders since 2003. It is clear that deposit interest rates collapsed after 2013, precisely the year when capital inflows pick up in Figure 4 and book values start increasing markedly in Figure 3.

Figure 5: Deposit Interest Rates, 2003–2024

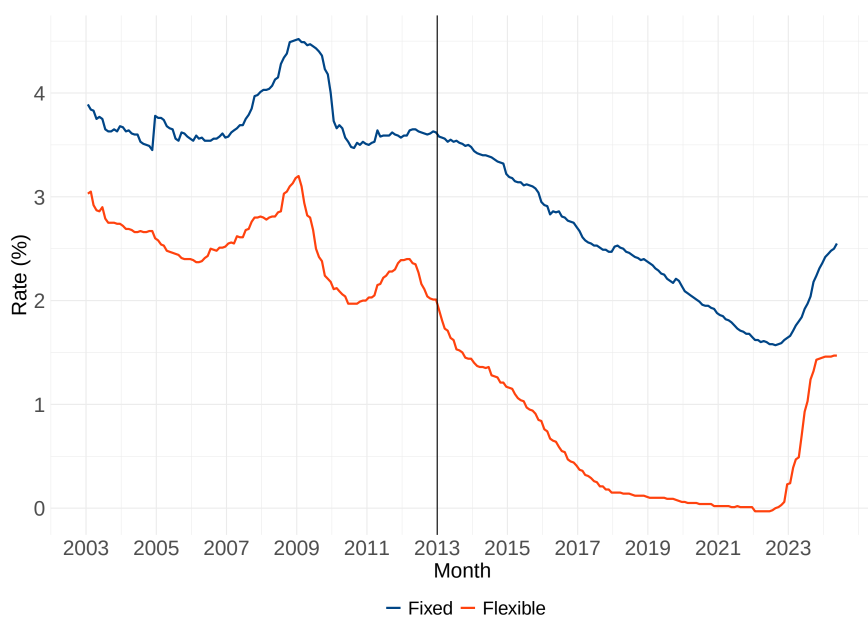

Notes: Figure shows monthly deposit interest rates for Dutch households. Red line shows flexible deposits, blue shows long-term fixed deposits. Data from De Nederlandsche Bank.

2013 is also the year in which some exemptions in Box 3 were phased out by the Rutte II government. These exemptions concerned social investments, cultural investments, and investments in venture capital. In particular the latter category may be relevant for the wealthiest investors. As those categories were no longer exempt from the wealth tax in Box 3 from 2013 onward, there were strengthened incentives to reallocate capital toward private firms in Box 2.

I conclude from the preceding analysis that the book values as recorded in the data are very sensitive to fiscal considerations. While precise causality cannot be established with my data, the timing of interest rate declines together with the ending of exemptions in Box 3 suggests that the increase in book values is to a significant extent artificial. In contrast, my alternative measures for firm value are based on firm capital and profits. While profits are in principle susceptible to fiscal manipulations, the corporate income tax environment remained quite stable in this period. Moreover, my use of three-year moving averages in $\overline{\Pi}_t$ should alleviate the worst of profit manipulations. Taken together, the fitted values V_t are more a more robust and plausible measure of firm value than the book values.

6.2 Parameter Implications

Another way of investigating the results from Section 5 is by pondering their implications for the model in Section 2. In standard neoclassical frameworks, a value of $\beta = 4.6$ would imply that on the margin, a euro of investment yields 4.6 euros in firm value. This is an implausibly high number; indeed, Tobin's average Q for U.S. nonfinancial corporations – which should equal marginal q under Hayashi (1982)'s assumptions – has not exceeded 1.75 since World War II, according to the U.S. Flow of Funds. In my model, there are no capital adjustment costs for simplicity; hence, absent market imperfections, Tobin's q should identically equal 1.

The point of the model is of course that market imperfections do exist and do matter. Equation (9) shows that β equals the product of two terms. The first term, $1/\mu$, is the inverse of the curvature of the profit function. μ depends on the capital output elasticity δ and the ratio of market power to returns to scale $\varphi = \lambda/\eta$. The second term, H_t, is a product of aggregate demand, the price level, and aggregate productivity, and also depends on δ and φ. The variables that enter H_t are unobservable to the econometrician, making it impossible to exactly disentangle the contributions of each variable to β. However, we can use the model to estimate all parameters separately, and investigate whether they are mutually consistent and make economic sense.

To identify the parameters, I follow Crouzet and Eberly (2023). In the model, Π_t is a firm's profits. Define υ as the ratio of profits to net revenue, and ϕ as the ratio of capital costs to operating surplus:

$$\upsilon := \frac{\Pi_t}{P_t C_t} \tag{30}$$

$$\phi := \frac{\rho K_t}{\Pi_t}. \tag{31}$$

Then, using equations (4)–(6), all reduced-form parameters can be identified, using

$$\varphi := \frac{\lambda}{\eta} = \frac{1}{\phi + (1-\phi)(1-\upsilon)} \tag{32}$$

$$\delta = 1 - \frac{1-\upsilon}{\phi + (1-\phi)(1-\upsilon)}, \tag{33}$$

$$\mu = \frac{\varphi - (1-\delta)}{\delta} = \frac{1}{\phi}. \tag{34}$$

To stay close to the model, I measure operating surplus as net revenue minus variable costs, where variable costs are the sum of total labor costs, material inputs, and bought services. The resulting measure for Π_t is very comparable to the series using net profits, but is generally a bit lower since profits also include the net extraordinary revenue (such as capital income from shareholdings). Net revenue is directly measured in the data, and is quite stable, with one exception. In 2017, measured net revenue jumps to about € 600 billion, more than doubling relative to surrounding years. Inspection of the data reveals that there are two firms in this year that each report revenue in excess of € 200 billion, while at the same time reporting extraordinary costs of also more than € 200 billion. Since all other firm-level variables for these firms are at much smaller magnitudes, this suggests that these firms engaged in fiscal manipulation and are dropped as outliers.[14] Capital costs are approximated using ρ^{wacc} found above, which should give the weighted-average cost of capital. I multiply ρ^{wacc} with the total value of the capital stock in a given year to approximate total capital costs.

Table 7: Parameters

Year	ϕ	υ	φ	δ	μ
2008	0.46	0.04	1.02	0.02	2.18
2009	0.50	0.05	1.03	0.03	2.02
2010	0.53	0.05	1.03	0.03	1.90
2011	0.92	0.03	1.00	0.02	1.09
2012	0.94	0.02	1.00	0.02	1.06
2013	1.02	0.03	1.00	0.03	0.98
2014	1.01	0.04	1.00	0.04	0.99
2015	0.69	0.05	1.02	0.04	1.46
2016	0.79	0.04	1.01	0.03	1.27
2017	0.82	0.04	1.01	0.04	1.22
2018	0.70	0.05	1.01	0.03	1.43
2019	0.59	0.06	1.02	0.03	1.70
2020	0.59	0.06	1.02	0.03	1.70
Average	**0.73**	**0.04**	**1.01**	**0.03**	**1.46**

Notes: ϕ is the ratio of capital costs to operating surplus; υ is the ratio of operating surplus to net revenue. φ is the reduced-form parameter governing returns to scale and markups; δ is the output elasticity of capital with respect to variable inputs; and μ gives the curvature of the profit function. All variables are aggregated each year to compute the parameters.

Table 7 reports the results. The ratio of capital costs to operating surplus, ϕ, is 0.73 on average. Over the years, this ratio fluctuates between 0.46 to 1.02 in 2013. Operating surplus is on average about 4% of total net revenue, which is quite stable. These two parameters combine to identify quite stable values for the markup/returns-to-scale ratio φ and the capital output elasticity δ. These values average at 1.01 and 0.03, respectively. Finally, the parameter we are most interested in, the profit curvature parameter μ averages at around 1.5.

The results imply that the second-order term in equation (9) is small but not negligible, since a $\mu \approx 1.5$ implies a coefficient premultiplying the second-order term of around -0.11. This suggests that omitting the quadratic term slightly biases the estimation results. The direction of this omitted-variable bias, however, is downward rather than upward, since the coefficient premultiplying the second-order term is negative, and the correlation between the first and second order terms is positive. My results should therefore be read as a lower bound.

These parameter estimates also help explain the large values found for β. Recall that β only equals Tobin's marginal q under neoclassical assumptions, which require no market power ($\lambda = 1$) and constant returns to scale ($\eta = 1$). As can be verified from Table 7, this condition is violated in almost every year, since the ratio of λ to η, $\varphi \neq 1$. The source of this violation

14. None of the other results in this paper are affected by these outliers, since these firms' capital stocks and profits are not outliers.

cannot be unpacked; Crouzet and Eberly (2023) show that λ and η are generally not separately identifiable. Nevertheless, the fact that $\varphi > 1$ indicates a positive net present value of rents. This positive value shows up in H_t and hence in β. This explains why, even absent any adjustment costs in capital, $\beta > 1$, and rationalizes why the coefficient is of a large magnitude (both β^- and β^+ are substantially larger than estimates of Tobin's q common in the literature).

The capital elasticity, δ, is surprisingly small. We are used to aggregate values of δ on the order of 0.3, not 0.03 (Barkai 2020). There can be two explanations for this divergence. First, in the model δ represents the elasticity with respect to all variable inputs. These inputs include labor but also materials and other variable costs. On the aggregate level, the production function is typically written using only labor and capital as inputs. There is in general no reason why micro and macro elasticities would have to be the same (Baqaee and Farhi 2019). A second explanation might be that the capital stock as recorded in firm balance sheets underestimates the capital stock at the aggregate level. Indeed, in Figure 3, the capital stock hovers around a value of €100 billion. From the National Accounts, we know that the aggregate capital stock (excluding housing) is much larger, hovering around €1 trillion in this period. Of course, not all of this tenfold increase can be allocated to firms' capital stocks; it is unclear, however, whether firms' capital stocks should be part of the corporate sector or the household sector. The household sector's capital stock is very close to the numbers reported in Figure 3 (Toussaint, de Vicq, Moatsos, and van der Valk 2022), suggesting that the capital stocks from firm balance sheets conform most easily to the household capital stock in the National Accounts.

Nevertheless, as discussed throughout the paper, capital stocks might be underrecorded both in the National Accounts and the microdata. Intangibles, in particular, are hard to measure and value. The upshot of mismeasured capital stocks for the parameters in Table 7 is the following. A larger K_t would increase the estimated costs of capital $\rho^{\text{wacc}} K_t$; hence, ϕ would increase one-for-one. The other parameters are affected by less than one-for-one. Moreover, not all the changes wrought by a larger K_t go in the "right" direction. Increasing K_t by a factor 10 – in effect assuming that firms own the entire capital stock – does result in a $\delta \approx 0.25$. However, the markups/returns-to-scale parameter φ declines to 0.8, which is less than 1 and hence not possible. Likewise, μ declines to 0.21, another impossibility. The reason for these countervailing changes is that while the capital stock might be increased, firm revenues and profits remain unchanged. Hence, to keep the model consistent, the other parameters have to decline.

We conclude that the model is useful for understanding aggregate dynamics for Dutch firms. I find higher curvature of the profit function, as measured by μ, than Crouzet and Eberly (2023). Even allowing for some mismeasurement in capital, this larger value is consistent with the structure of the firm accounts and hence is likely to be a robust finding. This suggests that private firms might have a larger flow value of rents than the public firms investigated in Crouzet and Eberly (2023). This rationalizes the large values for β found in the previous Section.[15]

7 Distributional Implications

7.1 Top Wealth Shares

Recall that the results in Table 5 converge to two coefficients: a high coefficient, β^+, and a low one, β^-. With these two values for β, I compute fitted values $\widehat{y}_{it}$, which under the identifying assumptions are free from measurement error. Hence, I will treat these fitted values as the true market values v_{it}. With these values of v in hand, we can update wealth shares. In Appendix B, I detail how I compute household's firm wealth. Throughout, I stay close to the official Statistics Netherlands procedure (Menger 2021). The brief procedure is as follows. We can link firms to their owners using the Shareholder Registry. I retain values for firms which can be linked to Dutch households. The Registry also provides information on ownership shares;

15. Crouzet and Eberly also investigate the properties of their model using National Accounts data from the United States. As discussed above, there is a general mismatch between aggregate statistics for the capital stock and microdata. Hence, if I were to use their model using Dutch National Accounts, I might well find different parameter values. For the purpose of this paper, however, the parameters reported in Table 7 confirm that the model is broadly consistent with my data.

hence, we can assign the updated values for firm wealth to households in proportion to their shareholdings. I deviate in one important aspect from Statistics Netherlands: Statistics Netherlands assigns ownership shares based on the share reported in the fourth quarter of the previous fiscal year, except for the first year in the dataset, where they use the first quarter of the current fiscal year. I adapt this procedure, by making the starting year firm-specific. If, for instance, firm A enters in 2008, that is the first year in my procedure, and I use the first-quarter ownership share for that year. Under the Statistics Netherlands procedure, this firm would not be included until 2009. As detailed in the Appendix, the Statistics Netherlands procedure assigns values for private business wealth to individuals who have no observable link to a firm, if they do declare firm dividend income. Since I have no meaningful improvement to those imputations, I keep those values, only replacing firm wealth for those households I can link to firms.

Once I have adjusted values of firm wealth at the household level, I recompute their position in the wealth distribution, where I adjust wealth totals as well. For instance, the series based on β^+ results in a new estimate of private business wealth. To compute top wealth shares for this series, I replace private business wealth for all matched households with this new measure, resum all new wealth totals, and recompute top wealth shares. This is important, since the adjusted values of private businesses could potentially accrue to households in the middle of the 'old' wealth distribution, who therefore by rights should be further up in the distribution than the microdata suggests.

This procedure results in three series for wealth shares: a raw series from the microdata, where the official aggregated business wealth is taken as given; and the fitted series based on β^+ and β^-. We begin by analyzing the top 1% wealth share. Figure 6 shows the unadjusted top 1% share in blue, as well as the two adjusted series. The high series, based on β^+, is in red; the low series based on β^- is in yellow.

Figure 6: Updated Top 1% Wealth Shares, 2008–2020

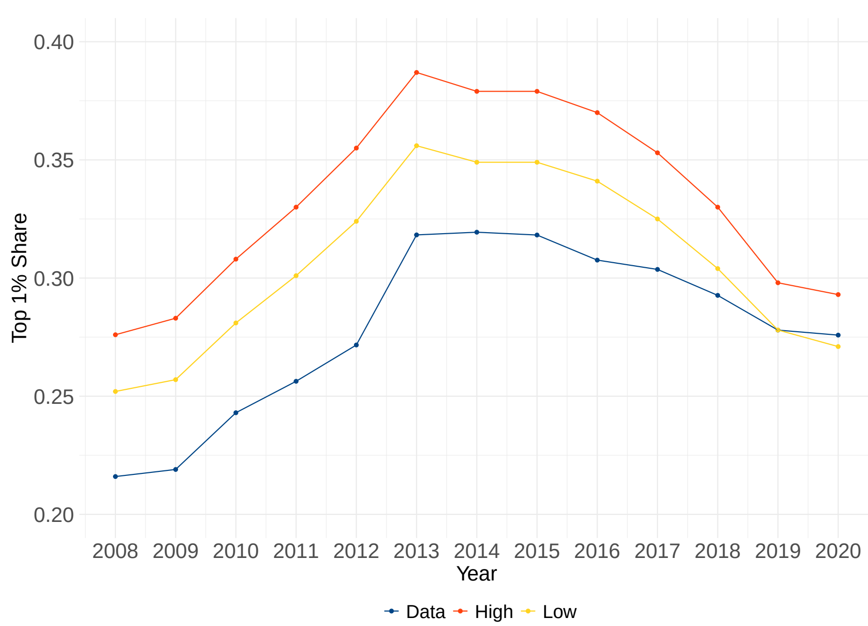

Figure 6 reports the results. In this timeframe, the microdata show a steep increase in wealth inequality, following the 2008 financial crisis, with the top 1% share rising from about 22% to almost one-third. After 2014 or so, inequality declined again, stabilizing at around 27% in 2020. The reasons for these dynamics are well-understood, and are due to differential portfolio composition of the richest households relative to the middle class (Toussaint, de Vicq, Moatsos, and van der Valk 2022; Toussaint 2022). The majority of households predominantly owns housing, whereas the rich hold most of their portfolio in private businesses and other financial assets. After the crisis, housing prices collapsed, while the stock market was faster to recover, resulting in a gain in wealth of the top 1% relative to most of the distribution. After 2014, housing prices started to increase rapidly, leading to a reversal of this trend. These dynamics mirror the "race between the housing market and the stock market" documented by Kuhn, Schularick, and Steins (2020).

The adjusted wealth shares also display these trends, but start on higher levels. Start with the high series, based on β^{+}. This series results in the highest wealth shares, which peak at 38.7% in 2013. The low series shows similar trends to the high series, but peaks at a lower level of 35.6% in 2013. This is still a substantial top wealth share, comparable to the United States (Saez and Zucman 2016; Smith, Zidar, and Zwick 2023).[16] The low series overlaps with the unadjusted top 1% share at the end of the period, stabilizing around 27.5%. The high series remains close to 30% even in those final years.

Both updated series rise and fall more or less similarly, showing that while the level of these series is determined by differences in the construction of the discount rate, their trends are driven more by time-series variation in firm profits. This nuances concerns by some authors that the trends in wealth inequality are primarily driven by interest rates (Greenwald, Leombroni, Lustig, and Van Nieuwerburgh 2021; Cochrane 2020). My results show that it is the interaction between discount rates and profits that matter; both prices and quantities, not just prices.

In conclusion, wealth shares are higher than recorded in the microdata. The extent of the upward adjustment differs based on specification. If we take the low coefficients as truth, the data should be adjusted by about 2.8 percentage points on average. If we take the high series as a reasonable average of the statistically strongest specifications, the adjustment should be in the order of 5.5 percentage points on average.

Next, we move to the top 0.1% share. Since private business wealth is even more concentrated among this bracket than among the broader top 1%, as revealed by Table 2, the resulting adjustments for the top 0.1% share should be even stronger. As before, Figure 7 plots the microdata together with the high and low series.

Figure 7: Updated Top 0.1% Wealth Shares, 2008–2020

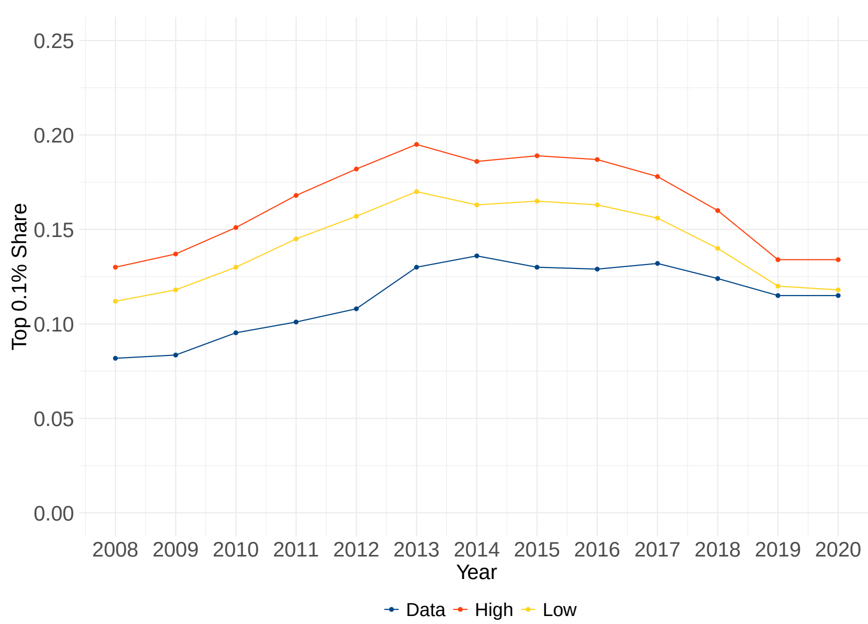

Interestingly, the microdata show a far less marked increase in top wealth shares through my sample, rising from about

16. For international comparisons, two complications arise. First, the unit of analysis differs between studies. In Saez and Zucman (2016) and Smith, Zidar, and Zwick (2023), the unit is a tax unit, which is not quite the same as a household. Second, the large values of Dutch private pensions may distort comparisons. As explained in Section 4, this wealth component is not included in Dutch wealth statistics, including this paper. Excluding it, however, ignores the large redistributive effect of the welfare state, which is less extensive in the United States. See Martínez-Toledano, Sodano, and Toussaint (2023) for further discussion.

7% to stabilize in the 11–12% range. Given the clear dynamics in the overall 1% share, this relative stability is puzzling. In contrast, both adjusted series show much stronger dynamics, which are fully in line with the dynamics observed for the top 1% share. The high series peaks at 19.5% in 2013, is stable around 19% until 2017, and declines in the final years to stabilize around 13.4% in 2020. The low series remains closer to the high series in relative terms than was the case for the top 1% share in Figure 6. The low series peaks at 17%, remains stable around 16.5%, then declines to 11.8%. For the high series, the upward adjustment clocks in at about five percentage points on average, which is similar in absolute terms to the upward adjustment in Figure 6, and therefore much more significant in relative terms. The low series shows an average upward adjustment of 2.8 percentage points, which is identical quantitatively to the adjustment in Figure 6. These results show that the extent of wealth concentration at the top is much larger than can be gleaned from the microdata, with the wealthiest households capturing most of the increase in firm value.

7.2 Wealth Composition

A second implication of the adjusted series is that the portfolio composition of the wealthiest may have changed. Figure 8 plots the old and adjusted shares of private business wealth in total wealth for the top 1%.

Figure 8: Adjusted Private Business Share of Total Wealth, Top 1%

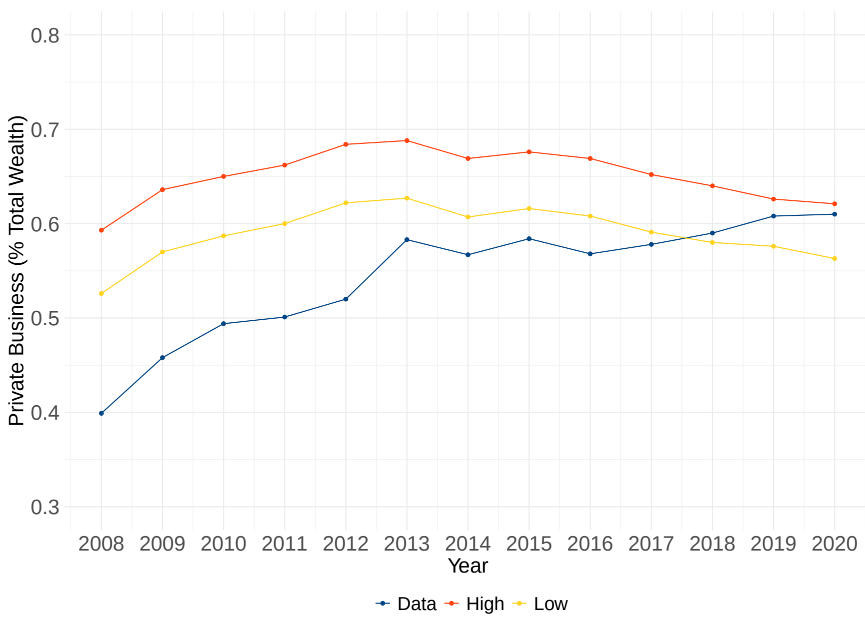

Figure 8 shows that private business wealth became increasingly dominant in the microdata. This could be both due to changes in prices and quantities. A decline in interest rates may have increased profitability and asset values of these corporations, boosting their value. At the same time, the richest may have increasingly reallocated their wealth towards private businesses. Fiscal considerations likely played a substantial role. As discussed in Section 4 and 6, households rebalanced financial assets toward private firms. Since firm-owners are concentrated at the top of the distribution, this portfolio reallocation results in a significant upward trend in the private firm share for the top 1%.

In contrast, the adjusted measures show a much more stable portfolio share. In these values, we also observe an increase between 2008 and 2013, but a relative stabilization afterward; and if anything, the private business share seems to decline somewhat after 2016. Since my alternative measures are based on firms' capital stocks and their profits, they are less sensitive to capital flows resulting from fiscal considerations. Thus, while the accounting importance of private business wealth for top portfolios has increased steeply in this period, the economic importance has remained stable.

8 Return Heterogeneity

In a final application, I turn to return heterogeneity. As is well known, standard incomplete-markets models have trouble generating the heavy tails of the wealth distribution seen in the data (Benhabib and Bisin 2018). Modifications that do match the levels of inequality require some multiplicative stochastic process, with some friction added to prevent wealth from becoming unboundedly negative after a succession of bad draws; see Gabaix (2009) for examples of such processes.[17]

In an important paper, Gabaix, Lasry, Lions, and Moll (2016) show that while such random-growth models can generate the levels of inequality, they cannot capture the speed of changes in inequality. To account for those dynamics, models need to be enriched such that returns correlate with wealth. The exact mechanism that generates this wealth-returns correlation is still unsettled (e.g., Jones and Kim 2018; Kacperczyk, Nosal, and Stevens 2019). Nevertheless, the centrality of return heterogeneity to accounting for inequality makes it a crucial concept to better understand.

Empirically, return heterogeneity has been documented in several instances. Piketty (2014) noted that U.S. universities with larger endowments earn higher returns; Saez and Zucman (2016) noted a similar scale effect for foundations. The most systematic evidence comes from Norway (Fagereng, Guiso, Malacrino, and Pistaferri 2020) and Sweden (Bach, Calvet, and Sodini 2020). These papers document wide heterogeneity in (realized or expected) returns across the wealth distribution. To a large extent, this is unsurprising, since portfolios vary dramatically and hence so do risk exposures (Cioffi 2021). However, Fagereng, Guiso, Malacrino, and Pistaferri (2020) show that even within narrow asset classes, returns increase with wealth.

Besides potentially accounting for inequality dynamics, this finding has major normative implications, since it provides a rationale for positive optimal capital (income) taxation (Gerritsen, Jacobs, Spiritus, and Rusu 2024; Guvenen et al. 2023; Guvenen, Kambourov, Kuruscu, and Ocampo-Diaz 2024). Hence, properly accounting for return heterogeneity is of first-order importance.

Unfortunately, these results are highly sensitive to even slight measurement error. This is particularly pressing for private business wealth, since most of the return heterogeneity documented in the literature is driven by the top of the distribution. As I have argued throughout the paper, private business wealth is subject to major measurement error. The problem with returns-wealth correlations is that it divides a cashflow (which potentially also has error) by wealth, which has measurement error, and then relates this ratio to wealth, which has the *same* measurement error. Hence, even absent true heterogeneity, left-hand and right-hand side would be mechanically correlated.

I formulate these claims in the following two propositions. Let $r^*_{it} := \pi_{it}/w^*_{it}$ be the (firm or firm-owner level) return, where π is the (correctly measured, for now) total cash-flow and capital gain out of true wealth w^*.[18] Now assume that wealth w^* is measured with error, so that we only observe $w = w^* + \xi$, with $\mathrm{E}[\pi\xi] = \mathrm{E}[w^*\xi] = 0$. Assume a finite variance of w^* equal to $\sigma^2_{w^*}$. Then, both returns and wealth will be mismeasured, and the common measurement error component will induce spurious correlation.

Proposition 1. *Consider the return $r^*_{it} := \pi_{it}/w^*_{it}$ where π is the correctly measured total cash-flow and capital gain out of wealth w^*. Then, when we consider the regression of returns on wealth $r = \beta w + \varepsilon$, the estimate $\widehat{\beta}$ is biased by*

$$\mathrm{E}\left[\widehat{\beta}\right] = \beta\left(1 - \frac{\sigma^2_{\xi}}{\sigma^2_{w^*} + \sigma^2_{\xi}}\right) - \mathrm{E}\left[r^*\xi^2\right]\mathrm{E}\left[w^*\right].$$

Proof. See Appendix A.3. □

17. Formally, these processes converge to a stationary Pareto distribution. Teulings and Toussaint (2023) show that top wealth is actually distributed Weibull, not Pareto. Teulings and Toussaint (2024) provide a network-based model with capacity constraints that can rationalize the Weibull distribution. Note that the test statistics developed in Teulings and Toussaint (2023) also show a clear rejection of Pareto in the dataset in this paper; results available upon request.

18. To emphasize that these regressions can run at the firm or household level, I use w^* instead of v to denote wealth.

The result is driven by the fact that a return is a ratio. The expectation of a ratio does not equal the ratio of expectations; instead, there are second-order terms which do not disappear. As a result, a returns regression is biased by two terms. The first term, $\left(1 - \frac{\sigma_\xi^2}{\sigma_{w^*}^2 + \sigma_\xi^2}\right)$ is the usual attenuation formula, which biases the coefficient toward zero. The second term, $-\mathrm{E}\left[r^*\xi^2\right]\mathrm{E}\left[w^*\right]$ is the result of the nonadditivity of the measurement error. The signs of these coefficients indicate that β will be attenuated toward zero.

Now include measurement error in cash-flows: $\pi_t := \pi^* + \xi$. This is plausible if capital gains are not properly measured, which is highly likely since market values are unobserved. This results in even worse bias in regressions, as formalized in the next proposition.

Proposition 2. *Include measurement error in cash-flows: $\pi_t := \pi^* + \xi$. Now, we have*

$$\mathrm{E}\left[\widehat{\beta}\right] = \frac{\mathrm{Cov}\left[r, w\right]}{\mathrm{Var}\left[w\right]} = \beta\left(1 - \frac{\sigma_\xi^2}{\sigma_{w^*}^2 + \sigma_\xi^2}\right) + \left(\sigma_\xi^2 - \mathrm{E}\left[r^*\xi^2\right]\right)\mathrm{E}\left[w^*\right]. \tag{35}$$

Proof. See Appendix A.3. □

Relative to Proposition 1, the second term is modified to $\left(\sigma_\xi^2 - \mathrm{E}\left[r^*\xi^2\right]\right)\mathrm{E}\left[w^*\right]$. The sign of this term is ambiguous and depends on the relative size of the variance of measurement error, σ_ξ^2, to the covariance between returns and squared measurement error. Hence, the bias in a regression where both numerator and denominator of returns are mismeasured is sizable but ambiguous in sign.

Note that if the measurement error in cashflows is not identical with that in wealth, but correlated, (i.e., $\pi = \pi^* + \zeta$, $\mathrm{E}\left[\xi\zeta\right] \neq 0$), the previous result goes through entirely, except that the σ_ξ^2 in the second term is replaced by $\mathrm{E}[\xi\zeta]$.

These two propositions underscore the importance of properly measuring private business wealth and returns. My GMM procedure is a first step toward addressing this issue. Since we have corrected values of firm value v and capital gains, bias will disappear from these regressions, and we can systematically investigate whether returns and wealth are correlated.

8.1 Measurement

I follow Fagereng, Guiso, Malacrino, and Pistaferri (2020) in constructing my returns measures. Details can be found in Appendix B. I consider returns both at the level of firm i and household j. Household j's return to firm i in year t is given by

$$r_{ijt} := \frac{\pi_{it} + \kappa_{it}}{v_{ijt-1} + \frac{1}{2}f_{it}} \cdot s_{ijt}. \tag{36}$$

Here, π are total profits, κ are capital gains, and s is the ownership share. As in Fagereng, Guiso, Malacrino, and Pistaferri (2020), the denominator not only includes beginning-of-period assets but also accounts for the fact that there are net inflows during the year, which might capitalize into profits and/or capital gains. Hence, the second factor corrects for the net inflows f, assuming that these occur about halfway during the years on average.

Aggregate this up to the household-level return as

$$r_{jt} := \sum_i r_{ijt} \cdot \underbrace{\frac{v_{ijt}s_{ijt}}{\sum_i v_{ijt}s_{ijt}}}_{=:\theta_{jt}}. \tag{37}$$

Hence, firm-level returns are aggregated to the household level using value weights θ_{jt}.

I calculate returns based on microdata for all firms and firm-owners in my sample, and do likewise with the adjusted firm-value series. For the microdata, I use total assets as my measure of firm value; this stays close to Fagereng, Guiso, Malacrino, and

Pistaferri (2020) and ensures that there are no interpretation difficulties when both numerator and denominator are negative, as is common at the bottom of the distribution. For my adjusted measures, I only have total value, not gross assets. However, since these values are mostly positive (since they are positively related to firms' capital stock, which is generally positive), this issue is less prevalent. I approximate κ in the microdata as the annual change in firm value net of in- and outflows of capital, which are separately recorded in my data and can thus be cleanly removed. For the adjusted series, I simply take the first difference of the adjusted values: $\kappa_{it} = \Delta v_{it}$. Since this corresponds purely to a change in value (since changes in profits are separately included in π_{it} and v_{it} is based on a smoothed three-year moving average of profits), this is a reasonable measure of capital gains.

For the results reported next, I drop all observations where returns do not exist (which happens if all components in Equation (36) do not exist or the denominator equals zero). Moreover, since small denominators can result in very large returns, I trim the top and bottom 0.5% of the sample in each year for all return measures. This results in a dataset of about 3.3 million firm-year observations.

8.2 Results

I report visual evidence for return heterogeneity in Figure 9. Panel 9a shows heterogeneity on the firm-level, with firms ranked by their position in the firm equity distribution. For each percentile of this distribution in each year, I collect all returns. For robustness, I then report the median of these returns for each percentile.[19] Panel 9b does the same for returns aggregated to the household level.

We observe clear return heterogeneity in all series. However, for the microdata, most heterogeneity is driven by a difference between the bottom 30% and the upper 70%. For the majority of the distribution, median returns are remarkably stable at around 8%. In fact, returns taper off at the top of the distribution, suggesting decreasing returns to scale in returns (as in Boar, Gorea, and Midrigan (2021)). These results also hold at the firm-owner level. At first glance, this seems to contradict results in Fagereng, Guiso, Malacrino, and Pistaferri (2020). Moreover, if high-value firms (or their owners) have *lower* returns than average, this would contradict most mechanisms proposed in the literature to tackle the Gabaix, Lasry, Lions, and Moll (2016) challenge.

Happily, the adjusted returns come to the rescue. Both series show a clear and steep gradient, suggesting that wealth is systematically correlated with returns. Moreover, the gradient increases above the 90th percentile, suggesting that the top of the distribution plays an outsized role in return heterogeneity. Moving from the median to the 100th percentile would cause a firm to see an increase in returns of almost 10 percentage points. For households, the gradient is less steep, but still sizable, with an increase in returns of more than 7 percentage points moving from the median to the top. The median returns using the adjusted values are now under the median returns using unadjusted data for most of the distribution. However, the gradient is much clearer than for the unadjusted data.

Figure 9 highlights the importance of properly measuring firm market values for return heterogeneity. However, median returns at each percentile might not be representative of overall return heterogeneity. To investigate these results more systematically, I run regressions of the form

$$r_{ijt} = \beta D_{ijt} + \varepsilon_{ijt}, \tag{38}$$

$$r_{jt} = \beta D_{jt} + \varepsilon_{jt}, \tag{39}$$

where D_{ijt} is a dummy for the decile of the distribution firm i is in. To make sure that the heterogeneity we capture is not

19. Results based on the mean for each percentile are broadly similar but with stronger volatility in the microdata series in the bottom 30%. This is presumably because even after the data cleaning, outliers remain due to denominator effects. Median estimates are robust to these outliers.

Figure 9: Return Heterogeneity over the Firm and Firm-Owner Distribution

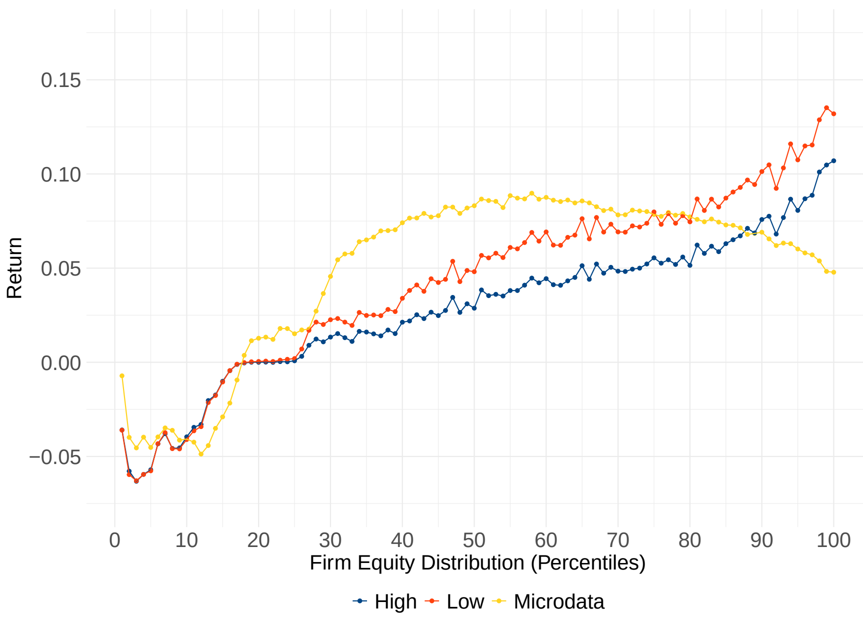

(a) Firms

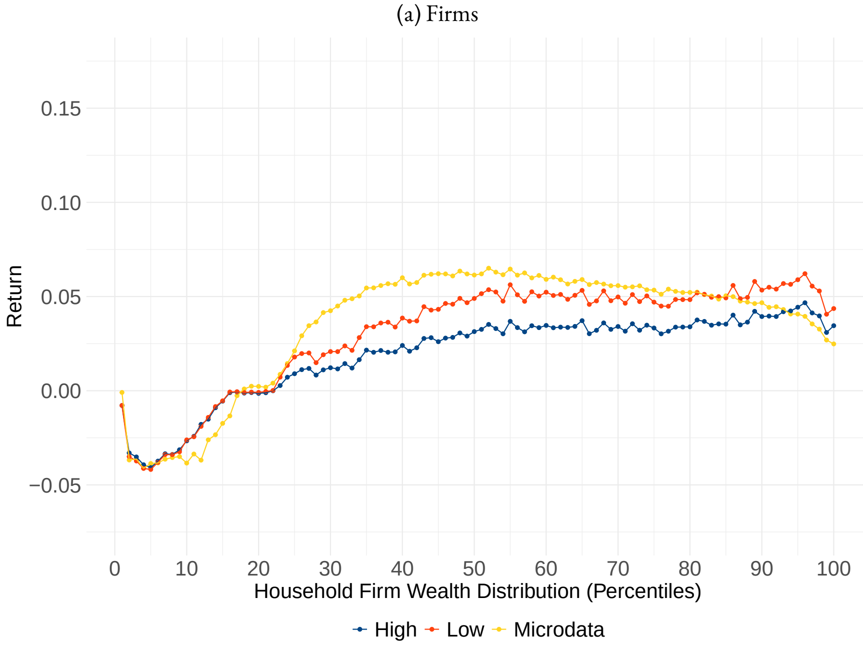

(b) Firm-Owners

Notes: Figures show heterogeneous returns along the distribution. Panel (a) shows the firm-size distribution, panel (b) the firm-owner distribution. Along each distribution in each year, firms are grouped in percentiles; per percentile (across years), the median return is calculated for the microdata series and the two adjusted series based on the WACC.

driven by systematic differences between firms, I also estimate these regressions with a battery of fixed effects. I successively investigate whether firm, year, and firm × year fixed effects affect the estimations. I run these regressions both at the firm level (equation (38) and the firm-owner level (39).[20]

20. Since a firm can be owned by multiple shareholders, there is a risk of double-counting in these regressions. However, Table 1 shows that the vast majority of firms are owned by a single owner, making this a minor concern.

Table 8: Return Heterogeneity Along the Firm Distribution

	r_{ijt}				r^+_{ijt}				r^-_{ijt}			
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	2.63				-0.060				-0.062			
	(1.82)				(0.110)				(0.133)			
Mean FE		8.536***	2.579***	1.042***		2.034***	-0.046	0.337***		2.635***	-0.051	0.438***
		(1.359)	(0.282)	(0.285)		(0.013)	(0.208)	(0.006)		(0.017)	(0.241)	(0.008)
Decile 2	-1.38	-2.96	-1.38	-0.020	0.120*	-0.270*	0.120*	0.283**	0.128	-0.422**	0.128	0.366**
	(1.33)	(2.67)	(1.33)	(0.012)	(0.052)	(0.111)	(0.052)	(0.083)	(0.063)	(0.131)	(0.063)	(0.091)
Decile 3	-0.677	-6.35	-0.677	-0.005	0.752***	-0.137	0.752***	0.696***	0.930***	-0.287	0.930***	0.888***
	(2.32)	(3.35)	(2.32)	(0.019)	(0.100)	(0.185)	(0.100)	(0.103)	(0.134)	(0.220)	(0.134)	(0.133)
Decile 4	0.451	-6.18	0.451	-0.022	1.28***	-0.031	1.28***	1.02***	1.61***	-0.133	1.61***	1.29***
	(2.59)	(3.42)	(2.59)	(0.015)	(0.112)	(0.237)	(0.112)	(0.153)	(0.150)	(0.284)	(0.150)	(0.201)
Decile 5	-2.12	-7.73*	-2.12	-0.016	1.59***	-0.008	1.59***	1.10***	2.01***	-0.090	2.01***	1.39***
	(2.27)	(3.49)	(2.27)	(0.014)	(0.133)	(0.264)	(0.133)	(0.188)	(0.176)	(0.320)	(0.176)	(0.243)
Decile 6	-2.23	-6.57	-2.23	-0.008	1.80***	-0.066	1.80***	1.32***	2.28***	-0.151	2.28***	1.68***
	(1.86)	(3.15)	(1.86)	(0.015)	(0.159)	(0.280)	(0.159)	(0.209)	(0.208)	(0.343)	(0.208)	(0.264)
Decile 7	-2.43	-6.73*	-2.43	-0.024	1.95***	-0.074	1.95***	1.55***	2.47***	-0.147	2.47***	1.97***
	(1.85)	(3.00)	(1.85)	(0.016)	(0.158)	(0.239)	(0.158)	(0.231)	(0.206)	(0.295)	(0.206)	(0.290)
Decile 8	-3.06	-6.93*	-3.06	-0.002	2.12***	-0.095	2.12***	1.71***	2.69***	-0.160	2.69***	2.17***
	(1.90)	(2.99)	(1.90)	(0.019)	(0.157)	(0.216)	(0.157)	(0.277)	(0.204)	(0.270)	(0.204)	(0.338)
Decile 9	-2.51	-6.58*	-2.51	-0.001	2.38***	-0.075	2.38***	2.02***	3.03***	-0.117	3.03***	2.56***
	(1.82)	(2.79)	(1.82)	(0.017)	(0.144)	(0.171)	(0.144)	(0.291)	(0.185)	(0.217)	(0.185)	(0.353)
Decile 10	-2.24	-4.90	-2.24	-0.002	2.76***	-0.145	2.76***	2.67***	3.51***	-0.191	3.51***	3.39***
	(1.92)	(2.92)	(1.92)	(0.021)	(0.171)	(0.153)	(0.171)	(0.365)	(0.219)	(0.190)	(0.219)	(0.436)
Fixed-effects												
Firm		✓				✓				✓		
Year			✓				✓				✓	
Firm × Year				✓				✓				✓
Fit statistics												
Observations	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312
R^2	0.000	0.561	0.000	1.00	0.009	0.363	0.014	0.991	0.009	0.363	0.013	0.991
Within R^2		0.000	0.000	0.000		0.000	0.009	0.012		0.000	0.009	0.012
RMSE	430.3	285.2	430.3	0.365	9.31	7.46	9.28	0.901	12.0	9.62	12.0	1.16

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Driscoll-Kraay standard errors in parentheses. r_{ijt} is the firm-level return based on unadjusted data, r^+_{ijt} is the firm-level return using β^+; r^-_{ijt} is the corresponding return using β^-.

Table 8 shows the results for the firm-level regressions. Because of multicollinearity, the first decile dummy drops out. I report the intercept (in the regressions without fixed effects) and the mean of the fixed effects, to indicate a benchmark against which the coefficients can be compared. I use Driscoll-Kraay standard errors to account for heteroskedasticity and autocorrelation.

The microdata-based returns (r_{ijt}) are insignificant in the first specification (column (1)), with none of the deciles having a systematically differential return. These conclusions are unchanged if we add fixed effects. Some decile dummies in column (2) are significant at the 5% level, indicating that using within-firm variation only, increasing firm size would *decrease* the firm return to wealth. Adding year or firm × year fixed effects again results in insignificant results. This indicates that accounting values cannot reveal systematic return heterogeneity using only within-firm variation. If anything, returns would seem to decline with scale, contradicting the return heterogeneity literature.

Contrast this with the results based on the adjusted returns. Across columns (5)–(12), results are highly significant and remain so after the inclusion of the fixed effects. Moreover, the gradient is sizable and increases in steepness at the top of the distribution. These results survive after the inclusion of the fixed effects. The results are qualitatively identical whether using r^{+}_{ijt} or r^{-}_{ijt}, although magnitudes differ somewhat. Results using only within-firm variation are generally insignificant (columns (6) and (10)). This is unsurprising: the adjusted returns are constructed at the firm level, leaving no room for within-firm variation. Within-year variation is significant, however. Moreover, using firm × year fixed effects does result in highly significant and economically sizable coefficients. This indicates that allowing firms to have separate time trends does aid identification.

My results indicate that return heterogeneity is sizable, mostly driven by the top, and is systematically related to firm wealth, even after allowing firms to have differential time trends and focusing on within-firm variation.

Table 9: Return Heterogeneity Along the Firm-Owner Distribution

	r_{jt}				r^{+}_{jt}				r^{-}_{jt}			
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Constant	2.54				-0.001				0.015			
	(1.36)				(0.096)				(0.118)			
Mean FE		6.115***	2.510***	0.634***		1.914***	0.010	0.827***		2.481***	0.022	1.068***
		(0.798)	(0.192)	(0.206)		(0.013)	(0.191)	(0.006)		(0.017)	(0.223)	(0.008)
Decile 2	-1.61	-3.59	-1.61	0.046**	0.160*	-0.176	0.160*	0.205*	0.182*	-0.291*	0.182*	0.255**
	(1.05)	(2.09)	(1.05)	(0.011)	(0.059)	(0.105)	(0.059)	(0.068)	(0.073)	(0.123)	(0.073)	(0.083)
Decile 3	-1.28	-4.75*	-1.28	0.080***	0.950***	-0.032	0.950***	0.445***	1.18***	-0.145	1.18***	0.563***
	(1.57)	(2.05)	(1.57)	(0.018)	(0.119)	(0.189)	(0.119)	(0.091)	(0.157)	(0.222)	(0.157)	(0.117)
Decile 4	-1.33	-5.54*	-1.33	0.090***	1.40***	0.032	1.40***	0.575***	1.77***	-0.041	1.77***	0.725***
	(1.86)	(2.40)	(1.86)	(0.019)	(0.131)	(0.224)	(0.131)	(0.076)	(0.170)	(0.268)	(0.170)	(0.100)
Decile 5	-2.28	-4.88*	-2.28	0.093***	1.61***	0.006	1.61***	0.642***	2.03***	-0.060	2.03***	0.809***
	(1.38)	(2.06)	(1.38)	(0.015)	(0.141)	(0.237)	(0.141)	(0.099)	(0.182)	(0.288)	(0.182)	(0.129)
Decile 6	-2.32	-4.74*	-2.32	0.097***	1.75***	-0.012	1.75***	0.767***	2.22***	-0.069	2.22***	0.968***
	(1.37)	(1.94)	(1.37)	(0.020)	(0.153)	(0.219)	(0.153)	(0.105)	(0.198)	(0.268)	(0.198)	(0.133)
Decile 7	-3.16	-5.22*	-3.16	0.101***	1.78***	-0.045	1.78***	0.892***	2.26***	-0.100	2.26***	1.13***
	(1.57)	(2.03)	(1.57)	(0.019)	(0.144)	(0.186)	(0.144)	(0.133)	(0.186)	(0.230)	(0.186)	(0.167)
Decile 8	-2.44	-4.76*	-2.44	0.103***	1.86***	-0.019	1.86***	1.03***	2.36***	-0.055	2.36***	1.31***
	(1.37)	(1.81)	(1.37)	(0.018)	(0.123)	(0.133)	(0.123)	(0.147)	(0.158)	(0.167)	(0.158)	(0.183)
Decile 9	-2.45	-4.39*	-2.45	0.104***	1.98***	-0.001	1.98***	1.14***	2.52***	-0.024	2.52***	1.44***
	(1.37)	(1.72)	(1.37)	(0.019)	(0.122)	(0.109)	(0.122)	(0.165)	(0.157)	(0.138)	(0.157)	(0.203)
Decile 10	-2.47	-2.96	-2.47	0.106***	1.97***	-0.040	1.97***	1.54***	2.50***	-0.067	2.50***	1.96***
	(1.37)	(1.98)	(1.37)	(0.022)	(0.122)	(0.089)	(0.122)	(0.209)	(0.158)	(0.109)	(0.158)	(0.257)
Fixed-effects												
Firm		✓				✓				✓		
Year			✓				✓				✓	
Firm × Year				✓				✓				✓
Fit statistics												
Observations	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312	2,898,312
R^2	0.000	0.481	0.000	1.00	0.006	0.364	0.011	0.987	0.006	0.364	0.010	0.987
Within R^2		0.000	0.000	0.001		0.000	0.006	0.005		0.000	0.006	0.005
RMSE	295.5	213.0	295.5	0.198	9.01	7.21	8.99	1.04	11.6	9.29	11.6	1.34

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*

Notes: Driscoll-Kraay standard errors in parentheses. r_{jt} is the firm-owner-level return based on unadjusted data, r^{+}_{jt} is the firm-owner-level return using β^{+}; r^{-}_{jt} is the corresponding return using β^{-}.

Table 9 does the same exercise for household-level returns. Results are mostly in line with those from Table 8, with some important differences. First, after including firm fixed effects, the accounting returns in column (2) are significantly negative for most of the deciles. Second, firm × year fixed effects now cause significant coefficients for the accounting returns in column (4), which are all positive. This indicates that allowing firms to have separate time trends results in household-level returns that have a positive gradient over the wealth distribution. In Table 9 as in Table 8, the adjusted return measures show a steeper gradient over the distribution than the accounting returns. The overall message is clear: return heterogeneity matters, mostly at the top, and particularly in the adjusted series.

These findings are not only informative for the return heterogeneity debate, they also help reconcile the firm-owner-level results of Fagereng, Guiso, Malacrino, and Pistaferri (2020) and others with the firm-level results by Boar, Gorea, and Midrigan (2021). These authors show that accounting returns exhibit decreasing returns to scale, and build a model to argue that financial constraints might drive a wedge between accounting returns and economic returns. My updated measures aim to capture the economic return to a firm, which is increasing in firm size. When aggregated to the household level, my adjusted returns continue to show a positive gradient. Accounting returns, on the other hand, are insignificant or even decreasing over the firm distribution (Table 8) but increasing when aggregated to the firm-owner level (Table 9). Hence, the results of Fagereng, Guiso, Malacrino, and Pistaferri (2020) and Boar, Gorea, and Midrigan (2021) are mutually consistent.

At the same time, my results also indicate that focusing on accounting returns can be misleading. When mapping theory to data, we are interested in the *economic* return: the compensation for risk-taking and management of the firm. My adjusted measures, while doubtlessly imperfect, are therefore closer to the corresponding variables in our economic models.

9 Conclusion

Private business wealth is a key determinant of distributional and aggregate dynamics. In this paper, I have developed a robust econometric procedure to identify the market value of private firms, based on time-series restrictions. The resulting series are consistent regardless of the choice of discount rate, result in more stable aggregate values of private business wealth, higher top wealth shares, and a steeper wealth-returns gradient.

My econometric procedure is straightforward to implement in settings where researchers can link firms to their owners. Even without this link, my results could be useful to adjust wealth shares. With data on firm balance sheets alone (such as from Compustat or Orbis), my GMM procedure could be implemented to create fitted values of private business wealth. Then, researchers could allocate the additional wealth along the wealth distribution in proportion to the weight of private business wealth in households' portfolios across the distribution. While imperfect, this adjustment is simple to do and would improve upon a status quo where accounting values are used as-is.[21]

My results have additional implications for the growing literature on return heterogeneity. Beyond its theoretical relevance, return heterogeneity has first-order effects for optimal taxation (Guvenen et al. 2023; Guvenen, Kambourov, Kuruscu, and Ocampo-Diaz 2024; Gerritsen, Jacobs, Spiritus, and Rusu 2024). My results show a steep returns-wealth gradient. This suggests a positive and sizable tax on wealth or capital income. Intriguingly, Guvenen, Kambourov, Kuruscu, and Ocampo-Diaz (2024) argue for a book-value tax as the optimal capital tax that minimizes economic distortions in the presence of return heterogeneity due to heterogeneous ability. Implementation of the book-value tax is robust to concerns of misspecification in my GMM procedure, but is theoretically justified by the return heterogeneity found using my procedure.

My results are a first step towards a systematic understanding of heterogeneous returns at the top of the distribution. Future work should consider mechanisms that generate return heterogeneity and wealth inequality and are consistent with the evidence found in this paper. Increasing availability of linked data, like in my setting, will hopefully permit econometric

21. Toussaint, van Bavel, Salverda, and Teulings (2020) implement such a procedure in the Dutch context; the top wealth shares found in that study are broadly similar to those found in this paper, suggesting that the simplified procedure is a reasonable quick fix in the absence of linked data.

evaluation of competing theories to explain the dynamics at the top of the wealth distribution.

References

Abel, Andrew B., and Janice C. Eberly. 1994. "A Unified Model of Investment under Uncertainty." *American Economic Review* 84 (5): 1369.

Anderson, T.W., and Cheng Hsiao. 1982. "Formulation and Estimation of Dynamic Models using Panel Data." *Journal of Econometrics* 18 (1): 47–82.

Andrews, Donald W.K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 59 (3): 817–858.

Andrews, Isaiah, James H. Stock, and Liyang Sun. 2019. "Weak Instruments in Instrumental Variables Regression: Theory and Practice." *Annual Review of Economics* 11:727–753.

Arellano, Manuel, and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies* 58 (2).

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2020. "Rich Pickings? Risk, Return, and Skill in Household Wealth." *American Economic Review* 110 (9): 2703–47.

Baqaee, David Rezza, and Emmanuel Farhi. 2019. "The Microeconomic Foundations of Aggregate Production Functions." *Journal of the European Economic Association* 17 (5): 1337–1392.

Barkai, Simcha. 2020. "Declining Labor and Capital Shares." *Journal of Finance* 75 (5): 2421–2463.

Benhabib, Jess, and Alberto Bisin. 2018. "Skewed Wealth Distributions: Theory and Empirics." *Journal of Economic Literature* 56 (4): 1261–91.

Bernstein, Asaf, and Peter Koudijs. 2024. "The Mortgage Piggy Bank: Building Wealth through Amortization." *Quarterly Journal of Economics* 139 (3): 1767–1825.

Blundell, Richard, and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics* 87 (1): 115–143.

Boar, Corina, Denis Gorea, and Virgiliu Midrigan. 2021. "Why Are Returns to Private Business Wealth So Dispersed?" *Working paper.*

Bruil, Arjan, Céline van Essen, Wouter Leenders, Arjan Lejour, Jan Möhlmann, and Simon Rabaté. 2022. "Inequality and Redistribution in the Netherlands." *CPB Discussion Paper* 436.

Campbell, Cole, and Jacob Robbins. 2023. "The Value of Private Business in the United States." *Available at SSRN 4635369.*

Cioffi, Riccardo A. 2021. "Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality." *Working paper.*

Cochrane, John. 2020. *Wealth and Taxes, part II.* Blog post.

Crouzet, Nicolas, and Janice C. Eberly. 2021. "Intangibles, Markups, and the Measurement of Productivity Growth." *Journal of Monetary Economics* 124:S92–S109.

———. 2023. "Rents and Intangible Capital: A Q+ framework." *Journal of Finance* 78 (4): 1873–1916.

Damodaran, Aswath. 2012. *Investment Valuation: Tools and Techniques for Determining the Value of Any Asset.* Vol. 666. John Wiley & Sons.

Erickson, Timothy, Colin Huan Jiang, and Toni M. Whited. 2014. "Minimum distance estimation of the errors-in-variables model using linear cumulant equations." *Journal of Econometrics* 183 (2): 211–221.

Erickson, Timothy, and Toni M. Whited. 2000. "Measurement Error and the Relationship between Investment and q." *Journal of Political Economy* 108 (5): 1027–1057.

———. 2012. "Treating Measurement Error in Tobin's q." *Review of Financial Studies* 25 (4): 1286–1329.

Fagereng, Andreas, Matthieu Gomez, Emilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik. 2024. "Asset-Price Redistribution." *World Inequality Lab Working Paper* 2024/14.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and Persistence in Returns to Wealth." *Econometrica* 88 (1): 115–170.

Fama, Eugene F, and Kenneth R French. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33:3–56.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Annual Review of Economics* 1 (1): 255–294.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The Dynamics of Inequality." *Econometrica* 84 (6): 2071–2111.

Gerritsen, Aart, Bas Jacobs, Kevin Spiritus, and Alexandra V. Rusu. 2024. "Optimal Taxation of Capital Income with Heterogeneous Rates of Return." *Economic Journal,* Forthcoming.

Gormsen, Niels Joachim, and Kilian Huber. 2023. "Corporate Discount Rates." *NBER Working Paper,* no. w31329.

———. 2024. "Firms' Perceived Cost of Capital." *NBER Working Paper* 32611.

Greenwald, Daniel L., Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh. 2021. "Financial and Total Wealth Inequality with Declining Interest Rates." *NBER Working Paper* 28613.

Griliches, Zvi, and Jerry A. Hausman. 1986. "Errors in Variables in Panel Data." *Journal of Econometrics* 31 (1): 93–118.

Gupta, Arpit, and Stijn Van Nieuwerburgh. 2021. "Valuing Private Equity Investments Strip By Strip." *Journal of Finance* 76 (6): 3255–3307.

Guvenen, Fatih, Gueorgui Kambourov, Burhan Kuruscu, Sergio Ocampo, and Daphne Chen. 2023. "Use It or Lose It: Efficiency and Redistributional Effects of Wealth Taxation." *Quarterly Journal of Economics* 138 (2): 835–894.

Guvenen, Fatih, Gueorgui Kambourov, Burhan Kuruscu, and Sergio Ocampo-Diaz. 2024. "Book-Value Wealth Taxation, Capital Income Taxation, and Innovation." *NBER Working Paper* 32585.

Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (4): 1029–1054.

Hayashi, Fumio. 1982. "Tobin's Marginal q and Average q: A Neoclassical Interpretation." *Econometrica,* 213–224.

IBO Vermogensverdeling. 2022. *Licht Uit, Spot Aan: De Vermogensverdeling.* Rijksoverheid.

Jacobs, Bas. 2015. "Belastingen op Kapitaalinkomen in Nederland." *Tijdschrift voor Openbare Financiën* 47 (1): 24–48.

Jones, Charles I., and Jihee Kim. 2018. “A Schumpeterian Model of Top Income Inequality.” *Journal of Political Economy* 126 (5).

Kacperczyk, Marcin, Jaromir Nosal, and Luminita Stevens. 2019. “Investor Sophistication and Capital Income Inequality.” *Journal of Monetary Economics* 107:18–31.

Kaplan, Steven N., and Antoinette Schoar. 2005. “Private Equity Performance: Returns, Persistence, and Capital Flows.” *Journal of Finance* 60 (4): 1791–1823.

Kapteyn, Arie, and Tom Wansbeek. 1983. “Identification in the Linear Errors in Variables Model.” *Econometrica* 51 (6): 1847–1849.

Karabarbounis, Loukas, and Brent Neiman. 2014. “The Global Decline of the Labor Share.” *Quarterly Journal of Economics* 129 (1).

Kopczuk, Wojciech, and Eric Zwick. 2020. “Business Incomes at the Top.” *Journal of Economic Perspectives* 34 (4): 27–51.

Korteweg, Arthur. 2019. “Risk Adjustment in Private Equity Returns.” *Annual Review of Financial Economics* 11 (1): 131–152.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins. 2020. “Income and Wealth Inequality in America, 1949–2016.” *Journal of Political Economy* 128 (9): 3469–3519.

Lewbel, Arthur, Susanne M. Schennach, and Linqi Zhang. 2023. “Identification of a Triangular Two Equation system Without Instruments.” *Journal of Business & Economic Statistics,* 1–12.

Martínez-Toledano, Clara, Alice Sodano, and Simon J. Toussaint. 2023. “Wealth Inequality in the Netherlands, 1894–2019: Updated Series With and Without Pension Wealth.” *WID. World Technical Note,* nos. 2023/13.

Menger, Jasper. 2021. “Herziening Statistieken Aanmerkelijk Belang vanaf 2006.” *CBS Publicatie.*

Montiel Olea, José Luis, and Carolin Pflueger. 2013. “A Robust Test for Weak Instruments.” *Journal of Business & Economic Statistics* 31 (3): 358–369.

Mood, Alexander McFarlane, Franklin A. Graybill, and Duane C. Boes. 1974. *Introduction to the Theory of Statistics.* McGraw-Hill Kogakusha.

Piketty, Thomas. 2014. *Capital in the Twenty-First Century.* Harvard University Press.

Reiersøl, Olav. 1950. “Identifiability of a Linear Relation between Variables Which Are Subject to Error.” *Econometrica* 18 (4): 375.

Saez, Emmanuel, and Gabriel Zucman. 2016. “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data.” *Quarterly Journal of Economics* 131 (2): 519–578.

———. 2020. “Trends in US Income and Wealth Inequality: Revising After the Revisionists.” *NBER Working Paper* 27921.

Schennach, Susanne M. 2016. “Recent Advances in the Measurement Error Literature.” *Annual Review of Economics* 8 (1): 341–377.

———. 2022. “Measurement Systems.” *Journal of Economic Literature* 60 (4): 1223–1263.

Smith, Matthew, Owen Zidar, and Eric Zwick. 2023. "Top Wealth in America: New Estimates Under Heterogeneous Returns." *Quarterly Journal of Economics* 138 (1): 515–573.

Teulings, Coen N., and Simon J. Toussaint. 2023. "Top Wealth Is Distributed Weibull, Not Pareto." *CEPR Discussion Paper* 18364.

———. 2024. "Why Has the Number of Billionaires Increased So Much?" *CEPR Discussion Paper* 19191.

Toussaint, Simon J. 2022. *Technisch Rapport Vermogensverdeling.*

Toussaint, Simon J., Bas J.P. van Bavel, Wiemer Salverda, and Coen N. Teulings. 2020. "Nederlandse vermogens schever verdeeld dan gedacht." *Economisch Statistische Berichten* 105 (4789): 438–441.

Toussaint, Simon J., Amaury de Vicq, Michail Moatsos, and Tim van der Valk. 2022. "Household Wealth and its Distribution in the Netherlands, 1854–2019." *World Inequality Lab Working Paper* 2022/19.

Xavier, Inês. 2020. "Wealth Inequality in the US: The Role of Heterogeneous Returns." *Working paper.*

A Mathematical Appendix

A.1 Closed-Form Expression of Profit Function

The derivation in this section follows Crouzet and Eberly (2023, Online Appendix IA.B.3.3) and is presented here for completeness. To economize on notation, define $\vartheta := (1-\delta)\eta$ and $\Psi_t := A_t K_t^{\delta\eta}$. Now the firm problem becomes

$$\begin{aligned} \Pi_t = & \max_{\{m_{gt}\}_{g=1}^{G}, P_t} P_t^{-\frac{\lambda}{\lambda-1}} D_t - \sum_{g=1}^{G} p_{gt} m_{gt}, \\ \text{s.t.} \quad & \Psi_t M_t^{\vartheta} \geq C_t, \end{aligned}$$

where $M_t := \prod_{g=1}^{G} m_{gt}^{\nu_g}$ and $C_t := P_t^{-\frac{\lambda}{\lambda-1}} D_t$. The dual to this maximization problem is to minimize variable costs:

$$\begin{aligned} \text{VC}_t = & \min_{\{m_{gt}\}_{g=1}^{G}} \sum_{g=1}^{G} p_{gt} m_{gt}, \\ \text{s.t.} \quad & \Psi_t M_t^{\vartheta} \geq C_t. \end{aligned}$$

The Cobb-Douglas structure of the problem ensures that the cost input share is equal to the output elasticity ν_g:

$$\frac{p_{gt} m_{gt}}{\mathcal{P}_t M_t} = \nu_g.$$

As a result, we have

$$\begin{aligned} \text{VC}_t &= \mathcal{P}_t M_t = \mathcal{P}_t \left(\frac{C_t}{\Psi_t}\right)^{\frac{1}{\vartheta}} \\ M_t &= \left(\frac{C_t}{\Psi_t}\right)^{\frac{1}{\vartheta}}. \end{aligned}$$

Substitution of these expressions into the original firm problem yields

$$\begin{aligned} \Pi_t = & \max_{\{m_{gt}\}_{g=1}^{G}, P_t} P_t^{-\frac{\lambda}{\lambda-1}} D_t - \mathcal{P}_t \left(\frac{C_t}{\Psi_t}\right)^{\frac{1}{\vartheta}}, \\ \text{s.t.} \quad & C_t \geq P_t^{-\frac{\lambda}{\lambda-1}} D_t, \end{aligned}$$

which has first-order conditions

$$\begin{aligned} P_t &= \lambda \text{MC}_t, \\ \text{MC}_t &= \frac{\mathcal{P}_t}{\vartheta C_t} \left(\frac{C_t}{\Psi_t}\right)^{\frac{1}{\vartheta}} \end{aligned}$$

so that:

$$C_t = \left(\frac{\lambda}{\vartheta}\right)^{-\frac{\vartheta\lambda}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} D_t^{\frac{\vartheta(\lambda-1)}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} \mathcal{P}_t^{-\frac{\vartheta\lambda}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} \Psi_t^{\frac{\lambda}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}}.$$

Inserting this into the solution for the cost-minimization problem, we get

$$\mathrm{VC}_t = \left(\frac{\lambda}{\vartheta}\right)^{-\frac{\lambda}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} D_t^{\frac{\lambda-1}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} \mathcal{P}_t^{-\frac{\lambda}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} \Psi_t^{\frac{1}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}}.$$

Profits are given by

$$\Pi_t = \left(\frac{\lambda}{\vartheta} - 1\right) \mathrm{VC}_t,$$

so that

$$\Pi_t = \left(\frac{\lambda}{\vartheta} - 1\right) \left(\frac{\lambda}{\vartheta}\right)^{-\frac{\lambda}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} D_t^{\frac{\lambda-1}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} \mathcal{P}_t^{-\frac{\lambda}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}} \Psi_t^{\frac{1}{\vartheta(\lambda-1)+(1-\vartheta)\lambda}}.$$

Using the definitions of ϑ and Ψ_t, these expressions can be rewritten as

$$\begin{aligned}
\Pi_t &= H_t^{1-\frac{1}{\mu}} K_t^{\frac{1}{\mu}}, \\
\mu &:= 1 + \frac{\varphi - 1}{\delta}, \\
\varphi &:= \frac{\lambda}{\eta}, \\
H_t &:= \left(\frac{\varphi}{1-\delta}\right)^{-\frac{\varphi}{\varphi-1}} \left(\frac{\varphi}{1-\delta} - 1\right)^{\frac{\varphi-(1-\delta)}{\varphi-1}} D_t^{\frac{\varphi-\eta}{\varphi-1}} \mathcal{P}_t^{-\frac{1-\delta}{\varphi-1}} A_t^{\frac{1}{\eta(\varphi-1)}}.
\end{aligned}$$

A.2 Derivation of Instrument Matrices

Let $\boldsymbol{R}_{\mathbf{FE}}$ be the within-transforming matrix, i.e.

$$\boldsymbol{R}_{\mathbf{FE}} = \boldsymbol{I} - \boldsymbol{J} = \begin{pmatrix} 1 & 0 & \dots & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & \dots & 1 \end{pmatrix} - \frac{1}{T}\begin{pmatrix} 1 & 1 & \dots & \dots & 1 \\ 1 & 1 & \dots & \dots & 1 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 1 & \dots & \dots & \dots & 1 \end{pmatrix}$$
$$= \begin{pmatrix} \frac{T-1}{T} & -\frac{1}{T} & \dots & \dots & -\frac{1}{T} \\ -\frac{1}{T} & \frac{T-1}{T} & -\frac{1}{T} & \dots & -\frac{1}{T} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ -\frac{1}{T} & \dots & \dots & \dots & \frac{T-1}{T} \end{pmatrix}.$$

It is easy to see that $\boldsymbol{R}'_{\mathbf{FE}}\boldsymbol{R}_{\mathbf{FE}} = \boldsymbol{R}_{\mathbf{FE}}$. Let $\boldsymbol{R}_{\mathbf{FD}}$ be the first-differencing matrix, i.e.,

$$\boldsymbol{R}_{\mathbf{FD}} = \begin{pmatrix} -1 & 1 & 0 & \dots & \dots & 0 \\ 0 & -1 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & \dots & -1 & 1 \\ 0 & \dots & \dots & \dots & \dots & -1 \end{pmatrix}.$$

Hence,

$$\boldsymbol{R}'_{\mathbf{FD}}\boldsymbol{R}_{\mathbf{FD}} = \begin{pmatrix} 1 & -1 & 0 & \dots & \dots & \dots & 0 \\ -1 & 2 & -1 & 0 & \dots & \dots & 0 \\ 0 & -1 & 2 & -1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & \dots & -1 & 2 & -1 \\ 0 & \dots & \dots & \dots & \dots & -1 & 1 \end{pmatrix}.$$

The second-difference matrix $\boldsymbol{R}_{\mathbf{SD}}$ is equal to $\boldsymbol{R}'_{\mathbf{FD}}\boldsymbol{R}_{\mathbf{FD}}$, hence:

$$\boldsymbol{R}'_{\mathbf{SD}}\boldsymbol{R}_{\mathbf{SD}} = \begin{pmatrix} 2 & -3 & 1 & 0 & \dots & \dots & \dots & \dots & 0 \\ -3 & 6 & -4 & 1 & 0 & \dots & \dots & \dots & 0 \\ 1 & -4 & 6 & -4 & 1 & 0 & \dots & \dots & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & \dots & 1 & -4 & 6 & -4 & 1 \\ 0 & \dots & \dots & \dots & \dots & 1 & -4 & 6 & -3 \\ 0 & \dots & \dots & \dots & \dots & \dots & 1 & -3 & 2 \end{pmatrix}.$$

The respective instrument matrices $\boldsymbol{P_1}$, $\boldsymbol{P_2}$ and $\boldsymbol{P_3}$ can be found as

$$\boldsymbol{P_1} = \boldsymbol{R}'_{\mathbf{FD}}\boldsymbol{R}_{\mathbf{FD}} - \boldsymbol{R}'_{\mathbf{FE}}\boldsymbol{R}_{\mathbf{FE}}, \tag{40}$$

$$\boldsymbol{P_2} = \boldsymbol{R}'_{\mathbf{SD}}\boldsymbol{R}_{\mathbf{SD}} - \boldsymbol{R}'_{\mathbf{FD}}\boldsymbol{R}_{\mathbf{FD}}, \tag{41}$$

$$\boldsymbol{P_3} = \boldsymbol{R}'_{\mathbf{SD}}\boldsymbol{R}_{\mathbf{SD}} - \boldsymbol{R}'_{\mathbf{FE}}\boldsymbol{R}_{\mathbf{FE}}. \tag{42}$$

$\boldsymbol{P_1}$ has the form:

$$\boldsymbol{P_1} = \begin{pmatrix} \frac{1}{T} & -\frac{T-1}{T} & \frac{1}{T} & \dots & \dots & \dots & \frac{1}{T} \\ -\frac{T-1}{T} & \frac{T+1}{T} & -\frac{T-1}{T} & \frac{1}{T} & \dots & \dots & \frac{1}{T} \\ \frac{1}{T} & -\frac{T-1}{T} & \frac{T+1}{T} & -\frac{T-1}{T} & \frac{1}{T} & \dots & \frac{1}{T} \\ \vdots & \frac{1}{T} & \ddots & \ddots & \ddots & \ddots & \vdots \\ \frac{1}{T} & \dots & \dots & \dots & \dots & -\frac{T-1}{T} & \frac{1}{T} \end{pmatrix}$$

and hence the vector of instruments $\boldsymbol{z_1}$ takes the form

$$\boldsymbol{z_1} \coloneqq \boldsymbol{Px} = \begin{pmatrix} \overline{x} - x_2 \\ \overline{x} - x_1 + x_2 - x_3 \\ \overline{x} - x_2 + x_3 - x_4 \\ \vdots \\ \overline{x} - x_{T-1} \end{pmatrix} \tag{43}$$

where $\overline{x} \coloneqq T^{-1} \sum_t x_{it}$ is the time-series average of x for firm i. This setup is easily extended to unbalanced panels by making the period length observation-specific, i.e., T_i instead of T.

$\boldsymbol{P_2}$ is of the form

$$\boldsymbol{P_2} = \begin{pmatrix} 1 & -2 & 1 & 0 & \dots & \dots & \dots & \dots & 0 \\ -2 & 4 & -3 & 1 & 0 & \dots & \dots & \dots & 0 \\ 1 & -3 & 4 & -3 & 1 & 0 & \dots & \dots & 0 \\ 0 & 1 & -3 & 4 & -3 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & \dots & \dots & 1 & -3 & 4 & -2 \\ 0 & \dots & \dots & \dots & \dots & \dots & 1 & -2 & 1 \end{pmatrix}$$

and hence we get an instrument vector $\boldsymbol{z_2}$ of the form

$$\boldsymbol{z_2} = \begin{pmatrix} x_1 - 2x_2 + x_3 \\ -2x_1 + 4x_2 - 3x_2 + x_4 \\ x_1 - 3x_2 + 4x_3 - 3x_4 + x_5 \\ \vdots \\ x_{T-3} - 3x_{T-2} + 4x_{T-1} - 2x_T \\ x_{T-2} - 2x_{T-1} + x_T \end{pmatrix}. \tag{44}$$

Finally, I construct an instrument $\boldsymbol{z_3}$, based on the difference between the within-estimator and the second-difference estimator. We get a matrix

$$P_3 = \begin{pmatrix} \frac{T+1}{T} & -\frac{3T-1}{T} & \frac{T+1}{T} & \frac{1}{T} & \cdots & \cdots & \frac{1}{T} \\ -\frac{3T-1}{T} & \frac{5T+1}{T} & -\frac{4T-1}{T} & \frac{1}{T} & \cdots & \cdots & \frac{1}{T} \\ \frac{1}{T} & -\frac{3T-1}{T} & \frac{5T+1}{T} & -\frac{4T-1}{T} & \frac{1}{T} & \cdots & \frac{1}{T} \\ \vdots & \frac{1}{T} & \ddots & \ddots & \ddots & \ddots & \vdots \\ \frac{1}{T} & \cdots & \cdots & \cdots & \cdots & -\frac{3T-1}{T} & \frac{T+1}{T} \end{pmatrix}$$

and hence an instrument vector

$$z_3 = \begin{pmatrix} \bar{x} + x_1 - 3x_2 + x_3 \\ \bar{x} - 3x_1 + 5x_2 - 4x_3 \\ \bar{x} - 3x_2 + 5x_3 - 4x_4 \\ \vdots \\ \bar{x} - 3x_{T-1} + x_T \end{pmatrix}. \tag{45}$$

A.3 Proofs of Propositions

Proof of Proposition 1. We have:

$$
\begin{aligned}
r &= \frac{\pi}{w} = \frac{\pi}{w^* + \xi} = r^* \left(1 - \frac{\xi}{w^* + \xi}\right), \\
\mathrm{E}[w] &= \mathrm{E}[w^* + \xi] = \mathrm{E}[w^*], \\
\mathrm{E}[r] &= \mathrm{E}[r^*] - \mathrm{E}\left[\frac{r^*\xi}{w^* + \xi}\right] \approx \mathrm{E}[r^*] - \frac{\mathrm{E}[r^*\xi]}{\mathrm{E}[w^* + \xi]} + \frac{\mathrm{Cov}[r^*\xi, w^* + \xi]}{\mathrm{E}[w^* + \xi]^2} - \frac{\mathrm{E}[r^*\xi]}{\mathrm{E}[w^* + \xi]^3}\mathrm{Var}[w^* + \xi] \\
&= \mathrm{E}[r^*] + \frac{\mathrm{E}\left[r^*\xi^2\right]}{\sigma^2_{w^*} + \sigma^2_{\xi}}, \\
\mathrm{E}[rw] &= \mathrm{E}[r^* w^*], \\
\mathrm{Cov}[r, w] &= \mathrm{E}[rw] - \mathrm{E}[r]\,\mathrm{E}[w] = \mathrm{E}[r^* w^*] - \left(\mathrm{E}[r^*] + \frac{\mathrm{E}\left[r^*\xi^2\right]}{\sigma^2_{w^*} + \sigma^2_{\xi}}\right)\mathrm{E}[w^*] \\
&= \mathrm{Cov}[r^*, w^*] - \frac{\mathrm{E}\left[r^*\xi^2\right]\mathrm{E}[w^*]}{\sigma^2_{w^*} + \sigma^2_{\xi}}, \\
\mathrm{Var}[w] &= \sigma^2_{w^*} + \sigma^2_{\xi},
\end{aligned}
$$

where we have repeatedly made use of the orthogonality of ξ with r^* and w^*, as well as the mean-zero expectation of ξ, and we have used a second-order approximation to the expectation of a ratio (Mood, Graybill, and Boes 1974). Collecting terms, we have

$$
\mathrm{E}\left[\widehat{\beta}\right] = \frac{\mathrm{Cov}[r, w]}{\mathrm{Var}[w]} = \frac{\mathrm{Cov}[r^*, w^*] - \frac{\mathrm{E}\left[r^*\xi^2\right]\mathrm{E}[w^*]}{\sigma^2_{w^*} + \sigma^2_{\xi}}}{\sigma^2_{w^*} + \sigma^2_{\xi}} = \beta\left(1 - \frac{\sigma^2_{\xi}}{\sigma^2_{w^*} + \sigma^2_{\xi}}\right) - \mathrm{E}\left[r^*\xi^2\right]\mathrm{E}[w^*].
$$

which is the result in the main text. □

Proof of Proposition 2. We follow the same steps as in Proposition 1. We have

$$\begin{aligned}
r &= \frac{y}{w} = \frac{\pi^* + \xi}{w^* + \xi} = r^* \left(1 - \frac{\xi}{w^* + \xi}\right) + \frac{\xi}{w^* + \xi}, \\
\mathrm{E}[w] &= \mathrm{E}[w^* + \xi] = \mathrm{E}[w^*], \\
\mathrm{E}[r] &= \mathrm{E}[r^*] - \mathrm{E}\left[\frac{r^*\xi}{w^* + \xi}\right] + \mathrm{E}\left[\frac{\xi}{w^* + \xi}\right] \\
&\approx \mathrm{E}[r^*] - \frac{\mathrm{E}[r^*\xi]}{\mathrm{E}[w^* + \xi]} + \frac{\mathrm{Cov}[r^*\xi, w^* + \xi]}{\mathrm{E}[w^* + \xi]^2} - \frac{\mathrm{E}[r^*\xi]}{\mathrm{E}[w^* + \xi]^3} \mathrm{Var}[w^* + \xi] \\
&+ \frac{\mathrm{E}[\xi]}{\mathrm{E}[w^* + \xi]} - \frac{\mathrm{Cov}[\xi, w^* + \xi]}{\mathrm{E}[w^* + \xi]^2} + \frac{\mathrm{E}[\xi]}{\mathrm{E}[w^* + \xi]^3} \mathrm{Var}[w^* + \xi] \\
&= \mathrm{E}[r^*] + \frac{\mathrm{E}\left[r^*\xi^2\right]}{\sigma_{w^*}^2 + \sigma_\xi^2} - \frac{\sigma_\xi^2}{\sigma_{w^*}^2 + \sigma_\xi^2}, \\
\mathrm{E}[rw] &= \mathrm{E}[r^* w^*], \\
\mathrm{Cov}[r, w] &= \mathrm{E}[rw] - \mathrm{E}[r]\,\mathrm{E}[w] = \mathrm{E}[r^* w^*] - \left(\mathrm{E}[r^*] + \frac{\mathrm{E}\left[r^*\xi^2\right]}{\sigma_{w^*}^2 + \sigma_\xi^2} - \frac{\sigma_\xi^2}{\sigma_{w^*}^2 + \sigma_\xi^2}\right) \mathrm{E}[w^*] \\
&= \mathrm{Cov}[r^*, w^*] - \frac{\left(\mathrm{E}\left[r^*\xi^2\right] + \sigma_\xi^2\right) \mathrm{E}[w^*]}{\sigma_{w^*}^2 + \sigma_\xi^2}, \\
\mathrm{Var}[w] &= \sigma_{w^*}^2 + \sigma_\xi^2,
\end{aligned}$$

hence, collecting terms, we have

$$\mathrm{E}\left[\widehat{\beta}\right] = \frac{\mathrm{Cov}[r, w]}{\mathrm{Var}[w]} = \beta \left(1 - \frac{\sigma_\xi^2}{\sigma_{w^*}^2 + \sigma_\xi^2}\right) + \left(\sigma_\xi^2 - \mathrm{E}\left[r^*\xi^2\right]\right) \mathrm{E}[w^*]. \tag{46}$$

□

B Data Construction

B.1 Overview

I use the following administrative datasets:

- VEHTAB: administrative records on the Dutch wealth distribution. All variables are dated January 1st of a year. VEHTAB contains anonymized household classifiers (`RINPERSOON`), as well as values for total wealth and all its components: total assets (consisting of deposits, other financial assets, private firm wealth, pass-through business equity, housing, other real estate, and other) and total liabilities (the sum of mortgages, student loans, and other loans).

- SZO AB+: administrative records linking firms to their owners, dated December 31st of a given year. SZO AB+ consists of two sources: Bedrijfsgegevens (BG), which contains the universe of corporate income tax records (merged with records from the Value-Added Tax, although these are less complete), and the Shareholder Registry (Aandeelhoudersregister, AR), which contains ownership information for all incorporated firms in each quarter of a fiscal year.

- BG contains, for each firm in each year, a full balance sheet and profit & loss statement. In addition, it includes information on firm age, number of employees, industry (at the five-digit level), and legal form. Note that not all variables are equally well observed. In particular, firm characteristics are not always observed. BG contains both consolidated and unconsolidated versions of balance-sheet and profit & loss variables.

- There are three kinds of identifiers: one household identifier (the `RINPERSOON`), and two types of firm identifiers:

  – BG has two variables for firm IDs per tax return: one for the firm filing the return (`RsinAangever_crypt`), and one for the firm for which the return is filed (`RsinAangegevene_crypt`). These two IDs may be the same, if the firm consists of one layer. For holdings and other multilayer structures, the uppermost layer (the "mother") will file for multiple firms, including itself and its "daughters".

  – AR has one variable for firm ID and one for ownership ID (which may be a firm or individual ID, i.e., a `RINPERSOON`). Hence, per firm-year observation, we have the firm ID and the ID of its owner. We also know if the owner is a corporation or natural person.

B.2 Linking Firms to Owners

I merge the BG and AR datasets for all years I have access to. I merge on the firm for which the return is filed, (`RsinAangegevene_crypt`). This creates a dataset which has all layers of all firms, and their ownership structure. We retain firms that are owned by natural persons, i.e., whose owner in the AR dataset is a `RINPERSOON`.

The firm in question may consist of a single layer, in which case there are no complications. Often, however, a firm has a complex structure, with a mother owning multiple daughters, which may in turn own firms. Mother firms report both consolidated and unconsolidated balance sheets and flow variables. We retain consolidated variables only, at the holding level. It may be that a household owns shares in a holding as well as direct shares in one or more of the daughters. If so, we retain the daughters as well, since by definition the holding will not own 100% of the daughter firms; hence, the mother balance sheets do not double-count.

This procedure will capture most of the holding structures that are linkable to individuals. However, complex ownership structures may still be undercounted. These structures are most often encountered at the top of the distribution (IBO Vermogensverdeling 2022). My results should therefore be seen as a lower bound on the true extent of inequality and heterogeneity of returns to wealth.

B.3 Estimating Private Business Wealth

Total wealth held by household j in firm i is simply obtained by multiplying j's share in the firm with i's wealth at beginning of the calendar year. Household variables are at January 1 and firm balance sheets at December 31, so I use lagged values of firm variables to obtain the relevant wealth. The data also has variables on beginning-of-year firm equity. Inspection reveals, however, that this variable is highly noisy, often not corresponding at all to the end-of-year values of one year earlier, where these should be the same. Hence, I opt to use lagged values of end-of-year firm equity throughout. The only exception is for the initial year 2007, where this is not possible. I follow the approach by Statistics Netherlands here (Menger 2021), which is to take the end-of-year balance sheet of 2007 but multiply it by ownership share in the first quarter.

My approach is similar to the one in Menger (2021), but there are some minor differences in sources used and aggregation rules. A major difference is that I impute missing values of ownership shares based on surrounding years: if at time t an ownership share is missing, but is present in either $t-1$ or $t+1$, that value is imputed (with the earlier date given preference). This results in more observations retained than in Menger (2021). Another difference is that sometimes the individual ID classifier is missing, which can similarly be imputed.[22]

The approach in Menger (2021) is as follows:

1. Determine net equity per firm, correcting for goodwill and other intangibles
2. Select ownership structures that generate wealth (excluding cooperations and so on)
3. Correct ownership shares:
   (a) Ownership shares are taken to be lagged shares on December 31 in year $t-1$, except for the first year 2007, where the share at the end of the first quarter of 2007 is used.
   (b) If ownership shares sum to more than 100%, they are proportionately scaled back.
   (c) If ownership shares are missing, there are two scenarios:
      i. The corporation is either listed, or large (>100 million in equity); in both cases, a conservative share of 5% is presumed
      ii. For all other scenarios, the capital is split equally among the observed shareholders
   (d) Remove outliers, which are defined as firms with at least a billion in equity but with no economic activity, which also have negative liquid assets (a potential sign of cooking the books).
4. Obtain indirect shares: this is important for a holding structure known as STAK (Stichting Administratiekantoor), which allows family offices to distribute shares among family members.
5. Impute missing observations: if a shareholder owns shares in a firm in $t-1$ and $t+1$ but ostensibly not in t, he is presumed to also own the firm in t.
6. Aggregate shares per person. At this stage, no correction is made for negative equity, even though these firms should be limited-liability. There are many fiscal constructs through which an owner can borrow from a firm, and as such the firm could have negative net worth.
7. If individuals receive firm dividends for at least three years, but have no observable claim to a firm, they are assigned firm wealth equal to their last received dividend multiplied by 10.

22. The anonymized individual IDs consist of a classifier followed by a unique string of digits. I match on the string of digits throughout. The classifier is either "R" (real person), or "F" (fictitious person, i.e., a corporation). If this classifier is missing altogether, CBS doesn't compute wealth totals, even though the digit-string and all other information is present, and even in cases where the classifier does appear in other years.

8. Aggregate to the household level. Here, negative equity is transformed to a symbolic value of 1 EUR.

My approach differs in step 3(a) and 3(c). Step 3(a) is problematic for firms whose first recorded year is later than 2007, since their entire first year is ignored. From a conceptual point of view, there is no good reason to treat those first years differently from the ones whose first year is observed in 2007. Hence, I apply the 2007 procedure for all first-year observations. Instead of fixed assignment rules like in 3(c), I impute ownership shares based on surrounding observations.

My data do not allow me to apply step 7, since I do not observe dividend payouts. In general, the private business wealth recorded in the wealth statistics VEHTAB have more observations than I am able to link to firms. For households which I cannot link to firms, I retain the private firm value recorded in VEHTAB. For households which I can link to firms, I replace recorded private business wealth with my measures estimated from my GMM procedure.

B.4 Constructing Returns to Wealth

Main variables the main variables I need for the main exercises are returns and gross assets. The return on assets of household i for firm j at time t is defined as

$$r_{ijt} := \frac{\pi_{jt} + \kappa_{jt}}{a_{ijt-1} + \frac{1}{2} f_{jt}} \cdot s_{ijt}. \tag{47}$$

here, π are total profits, κ are capital gains, and s is the ownership share. As in Fagereng, Guiso, Malacrino, and Pistaferri (2020), the denominator not only includes beginning-of-period assets but also accounts for the fact that there are net inflows during the year, which might capitalize into profits and/or capital gains. Hence, the second factor corrects for the net inflows f, assuming that these occur about halfway during the years on average. In what follows, t denotes a calendar year, where all flow variables are occurring during the year, and all stock variables are dated at December 31 of a given year. I match all of these objects to my data as follows:

- π_{jt} = `OND1_0789_BELASTBARE_WINST`, which is the total taxable profit of firm j at time t
- κ_{jt} = `OND1_0793_TOTAAL_VERMOGENSVERSCHIL`, which is the net difference in total wealth between $t-1$ and t, net of inflows and outflows of capital.
- $a_{ijt-1} = a_{jt-1} \cdot s_{ijt-1}$ = `OND1_1376_TOTAAL_ACTIVA_EB * BetrokkenheidUltKw4`, total assets at the end of calendar year $t-1$, multiplied by the share at the end of the fourth quarter of period $t-1$.
- f_{jt} = `OND1_0803_STORTINGEN_VAN_KAPITAAL_IN_HET_BOEKJAAR` – `OND1_0798_TERUGBETALING_KAPITAAL_PRIVE_ONTTREKKING`, i.e., the difference between inflows of capital during the year and outflows (either repayments or withdrawals for private use). The difference is therefore the net inflow of capital during period t.
- s_{ijt} = `BetrokkenheidUltKw4`, the ownership share at the end of the fourth quarter of period t.

As discussed in the main text, for the updated returns measures, I replace assets a_{jt} with updated firm wealth v_{jt} and measure capital gains κ as the first-difference of updated firm wealth, $\kappa = \Delta v_{jt}$.

C Additional Figures and Tables

The GMM estimations in the main text use heteroskedasticity-robust standard errors. If we instead cluster standard errors at the firm level, we get the estimations in Table 10.

Table 10: GMM Estimation, Time-Series Identification, Clustered Standard Errors

Dependent Variable:	y_{it}						
Instruments	z_1	z_2	z_3	z_1, z_2	z_2, z_3	z_1, z_3	z_1, z_2, z_3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: ρ^{wacc}							
β	4.610**	4.089*	4.543	4.040*	4.224**	4.597**	2.859***
	(1.704)	(1.819)	(2.415)	(1.831)	(1.553)	(1.665)	(0.711)
Observations	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563
J-test Statistic				0.870	0.027	0.001	2.611
J-test p-value				0.351	0.870	0.975	0.271
First Stage F-statistic	10.889	10.095	10.486	9.131	7.367	11.392	9.050
Panel B: ρ^{gh}							
β	3.242**	2.909*	3.580	3.061*	3.061**	3.237**	2.831**
	(1.169)	(1.298)	(1.959)	(1.223)	(1.117)	(1.164)	(0.874)
Observations	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563
J-test Statistic				0.505	0.086	0.037	1.887
J-test p-value				0.478	0.769	0.848	0.389
First Stage F-statistic	10.889	10.095	10.486	9.131	8.951	11.392	9.050
Panel C: ρ^{b}							
β	6.347**	5.289*	5.052*	4.706*	5.179**	5.570*	2.645***
	(2.252)	(2.350)	(2.256)	(2.325)	(1.965)	(2.243)	(0.469)
Observations	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563	6,713,563
J-test Statistic				1.668	0.006	0.307	4.417
J-test p-value				0.196	0.936	0.579	0.110
First Stage F-statistic	10.889	10.095	10.486	9.131	8.951	11.392	9.050

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05*
Notes: Two-step GMM with standard errors clustered at the firm level in parentheses. Robust first-stage F-statistic by Montiel Olea and Pflueger (2013). Each panel shows GMM results with as independent variable the $\boldsymbol{y}$ computed using the respective discount rate.

The results are broadly comparable, with the exception of column (7), which is uniformly lower across specifications (and is almost identical across panels). Since heteroskedasticity is a major concern in my setting (see Table 1), I stick with the specifications in the main text.